Side Channels, Compromising Emanations and Surveillance: Current and future technologies

Richard Frankland

Technical Report
RHUL–MA–2011–07
8th March 2011

Department of Mathematics
Royal Holloway, University of London
Egham, Surrey TW20 0EX, England

http://www.rhul.ac.uk/mathematics/techreports

Side Channels, Compromising Emanations and Surveillance: Current and future technologies.

by

Richard Frankland

Student Number: 100500485

Supervisor: Prof. Keith Martin

Submitted in partial fulfillment of

the requirements for the degree of

Master of Science

in

Information Security

at the

Department of Mathematics

Royal Holloway, University of London

I declare that this assignment is all my own work and that I have acknowledged all quotations from the published or unpublished works of other people. I declare that I have also read the statements on plagiarism in Section 1 of the Regulations Governing Examination and Assessment Offences and in accordance with it I submit this project report as my own work.

Signature _____ Date _____

Acknowledgments

I would like to thank my project supervisor, Prof. Keith Martin, for his incredibly helpful feedback and guidance, which have greatly aided the development of this project.

I would also like to thank my parents for their encouragement and support during these past months, which have been an invaluable source of inspiration to me.

Abstract

Side channel attacks exploit implementation-specific information leakage to defeat cryptographic measures, usually designed to provide confidentiality. The majority of published attacks require physical possession or network access to the target device, and as such are not suitable for implementation by an attacker wishing to carry out a remote, passive and undetectable attack. Surveillance techniques can be applied to compromising emanations, which is effectively side channel leakage from devices handling the sensitive, and protected, information desired by the attacker. By monitoring compromising emanations from devices processing data in plaintext, such as computer monitors and keyboards, an attacker can completely bypass the protection offered by cryptographic primitives. Published work on these emanations focusses on three main sources; electromagnetic radiation, visual representations of the data, and acoustic signals. This project aims to review the literature on techniques exploiting these emanations, and place them in the context of real world attacks, potentially undertaken by a variety of individuals and organisations. Current and future technological developments of sensors and consumer electronics will also be discussed in relation to their applicability to these attacks, identifying possible directions in which they might develop in the future.

Contents

List of Figures

Chapter 1

Introduction

Recovery of plaintext data is the goal of any attacker seeking to break the confidentiality provided by cryptography. Practical implementations of logical and side channel cryptanalysis both seek to do this by exploiting faults in either the design of the cipher or its implementation respectively, thus recovering the secret key or the plaintext. However, attempting to do so in a real world setting may require specific, restrictive, conditions to be achieved, such as network access or physical proximity to the protected device. Even then, evidence of an attack may be apparent, something which could be undesirable.

In such a case, surveillance techniques would seem to offer an ideal solution. By purely monitoring the use of the protected device an attacker could gain information required to access it, such as a password, or even the protected data itself, perhaps through a visual representation. Furthermore, by looking at design faults that compromise confidentiality in the device itself, one could interpret processes to extract leaked information as side channel analysis of the device. In fact, these sources of leaked information can be considered as compromising emanations, leaking plaintext data that can be recovered for analysis through a passive, remote and undetectable attack. Most interestingly, these attacks can be used to bypass strong cryptographic protection mechanisms; a password for an encrypted hard drive can be eavesdropped or an encrypted email or document can be reconstructed just from emanations while being typed or displayed on screen.

The first aim of this project is to put attacks utilising compromising emanations into context with real world implementations of surveillance technology by investigating and analysing the nature of these attacks through the review of available literature. Their current status and practicality will be assessed, as will their applicability in the current security landscape, and possible implementations using current and future technologies will also be suggested and discussed. The second aim of this project is to highlight how the feasibility of such attacks, with the potential for further development through new research in fields such as signal processing and mobile computing, and the increasing implementation of cryptographic security measures may result in instances of recorded attacks in the wild, with plaintext data leakage being an attractive target as a weak link in the chain of security.

Beginning with the concept of the surveillance of side channel leakage and the exploitation of compromising emanations for real world attacks, this first chapter covers the development of cryptographic side channel analysis and identifies its limitations with regards to application as part of surveillance attacks on confidentiality, with the suitability of attacks focussing on plaintext data leakage later being discussed. The main portion of the project is made up of reviews of published attacks which make use of compromising emanations. These reviews are split into chapters by their shared source of emanations and ordered chronologically from the year of first published attack;

- Chapter 2 explores the development of attacks for recovering plaintext data from electromagnetic emanations and ends with a short analysis of techniques and technology used for attack and defence.
- Chapter 3 examines the implementation of novel optical attacks, and posits the application of new imaging and display technology for attack and defence methods.
- Chapter 4 covers the entire development of techniques exploiting acoustic emanations, and suggests possible methods of attack and defence.
- Chapter 5 addresses the aims of this project, discussing them in relation to the information reported in the above chapters. Possible avenues of future research and technological development are also identified, and their relevance examined.
- Chapter 6 concludes the report, reiterating the power of the attacks, namely their ability to recover plaintext information, potentially breaking the confidentiality offered by cryptographic implementations. Finally, it is surmised that the continued development of technology that can be used for attack and defence, and the continued open publication of methods utilising compromising emanations, will ensure the importance and relevance of these types of attack for the foreseeable future.

1.1 Definition of terms

Before going further into the reason for reviewing the application of side channel analysis for surveillance purposes, and their potential for the development of new and powerful attack methodologies, it is worth defining what exactly these two terms mean, and their relevance to this project.

1.1.1 Definition of side channel

The term “side channel” refers to an implementation-specific form of information leakage, usually from a cryptographic implementation, in a manner not considered in the data flow model of the implementation [73]. This leakage can allow the determination of the secret key, and reconstruction of plaintext data, breaking the confidentiality offered by a cryptographic primitive. It is important to note that a successful side channel attack does not equate to a successful break of the cryptographic algorithm, only demonstrating that the method of implementation, which can be in software or hardware, is not secure.

1.1.2 Definition of surveillance

With regards to UK law, "surveillance" can be defined as the monitoring, observing, and listening to persons, their movements, conversations, other activities or communications and the recording of anything monitored, observed or listened to in the course of surveillance, by or with, assistance of a surveillance device [67]. It can be split into two main types, directed, undertaken with prior planning, and intrusive, involving the physical presence of surveillance technology or personnel on residential premises or in a vehicle, both being carried out covertly. Mass surveillance is a term commonly used to describe surveillance infrastructure like CCTV or Automatic Number Plate Recognition systems, enacting a kind of "dragnet" approach, starkly contrasting with targeted surveillance methods.

Within the scope of this project, surveillance will be considered, theoretically, as any covert monitoring of a specific target, of which information may already be known, especially technological capabilities. The purpose of this monitoring is to recover any sensitive information that may or may not be protected by the target using modern data protection methods. Since it can be considered a remote, passive attack, bypassing the confidentiality provided by data protection measures, in this project any implementer of the outlined attack methods will be called the "attacker", while the subject against whom the attacks are directed will be called the "target".

By combining techniques for side channel exploitation with existing surveillance methodology, information leakage from devices being used by a target can be utilised by an attacker to defeat any measures designed to provide confidentiality. The development of such attacks can be used as the basis of more powerful surveillance techniques for future implementation.

1.2 A short introduction to side channels

The development of side channel attacks originally came from methods of breaking the confidentiality provided by cryptographic primitives by using implementation-specific information leakage. In this vein, the first side channel attacks exploited differences in the length of time it took for certain operations of a cryptographic algorithm to complete.

1.2.1 Timing attacks

Timing attacks first came into prominence with Kocher's work on public key algorithms that relied on modular exponentiation [48]. The general attack relies on the ability of an attacker to eavesdrop on the known cryptographic protocol and collect time measurements of secret exponent operations on several known plaintexts. Individual bits of the key can be guessed with high probability through comparison of time variations that arise due to differences in the speed of modular exponentiation calculations, depending on accuracy of timing measurements. The method was demonstrated practically on a deprecated version of a "CASCADE" smart card running RSA using the Montgomery algorithm [32]. Improvements on the attack also allowed

successful key recovery from implementations of RSA using the Chinese remainder theorem, by using time differences to calculate an integer multiple of either p or q, the distinct primes that are used to generate the both the RSA modulus and the secret exponent [75].

Symmetric algorithms have also been the target of timing attacks, again focussing on discrepancies in time taken for specific operations. Bad implementations of Rijndael have been shown to be vulnerable to timing attacks, where certain operations do not run in constant time [50], and also across networks in certain versions of OpenSSL implementing AES [15], where key bits were found to be recoverable over the network itself using timing information from the computer performing the encryption operations. Implementations of RSA in OpenSSL have also proved vulnerable to remote timing attacks across networks [19].

Since the most effective countermeasures against these attacks involves preventing leakage through system design, making sure operations are performed in constant time, the majority of modern cryptographic software is now resistant to timing attacks. However, it should be noted that most side channel methods have a timing component to them.

1.2.2 Power and electromagnetic analysis

Power analysis attacks evolved from the work done on timing attacks, with a greater emphasis on attacking hardware implementations. Instead of just timing intervals, differences in power consumption during encryption operations, measured over time, using power probes in conjunction with oscilloscopes, are utilised to determine information about the secret key [49]. In that paper, two main types of power analysis were developed; “SPA”, simple power analysis, and “DPA”, differential power analysis. SPA of DES, for instance, was shown to produce a power trace that can elucidate the information about a single encryption operation, while it was demonstrated how DPA of DES can make use of multiple traces to determine individual key bits by multiple guesses. Further power analysis attacks have been carried out against implementations of other cryptographic primitives; the stream ciphers Grain and Trivium [38], an HMAC based on SHA-256 [63], and also on AES [18].

Electromagnetic (EM) analysis of cryptographic implementations derives from the techniques used in power analysis attacks, with “SEMA”, simple electromagnetic analysis, and “DEMA”, differential electromagnetic analysis, existing as EM counterparts to the different types of power analysis. This analogy was first suggested in Kocher et al.’s original paper on power analysis, indicating that EM analysis was, at that stage, already being seen as the next logical progression for practical side channel attacks [49].

The main difference between EM analysis and power analysis is that in an EM attack power consumption is measured through the detection of electromagnetic currents caused by electric activity in the target device over time, the strength of the current also being measured, allowing more information to be inferred about the cryptographic process [71]. The attack has been demonstrated against implementations of DES and RSA in smart cards, with successful extraction of secret keys from versions not using software countermeasures [40]. In that experiment, an EM probe, a hand made copper solenoid, was placed in close proximity to the tested

smart card, as close to where the CPU was as possible. EM traces were collected and used to determine key bits in much the same method as power analysis techniques.

Further development of EM analysis lead to more practical attacks; in a more realistic scenario, EM attacks, using various differential methods, on implementations of Rijndael and elliptic curve cryptography written in Java, have been demonstrated on a PDA, resulting in recovery of the secret key [41]. The practicality of EM attacks is receiving greater attention through research into low-cost implementable attacks against newer technologies using established cryptographic primitives, such as contactless smart cards implementing the banking industry standard, 3DES [45].

The techniques for both power analysis and EM analysis, described above, focus on defeating cryptographic implementations, usually cryptographic primitives forming part of a suite implemented in software, such as OpenSSL, or in embedded hardware as in a smart card. The goal of such an attack is generally recovery of the secret key, which can be found using side channel information gained by having intimate access to the target device. This requires physical adjacency, with probes and oscilloscopes, and also possession of the device being attacked. Despite the demonstrable power of these attacks, from a surveillance perspective; remote, passive and undetectable attack techniques are required for the covert monitoring and recording of sensitive information, something that cannot be achieved by using existing side channel methodologies.

1.3 Compromising emanations: plaintext information leakage and surveillance

There is one ubiquitous aspect of cryptographic implementations that provide confidentiality between two users; the data being protected must at some point be human readable. Whether it is at the input stage, or when the plaintext is presented on a computer display or some other readable format, sensitive data is vulnerable to eavesdropping before the encryption process and after the decryption process. Here, links to traditional surveillance can come into play, and can be directed against common methods of data input and display. By targeting such devices as computer keyboards and monitors it is possible for one to recover information through remote, passive, observation. The most obvious way of doing so would be by visual monitoring, performed by a surveillance operative and standard optical or imaging equipment, such as a telescope and camera. However, by investigating and exploiting device-specific information leakage there may be other ways to accurately, and autonomously, recover sensitive data.

For a passive attack, the source of this leakage must come from side channel data being given out by the target device, and can be thought of as a “compromising emanation”. These emanations can take the form of any distinct signal containing leaked information from a device, but in the academic literature three distinct types are the most published, and have the greatest relevance to known surveillance techniques. These are;

- Electromagnetic emanations, taking the form of EM radiation in the radio frequency spectrum leaking plaintext data

- Optical emanations, leakage through visual representation of sensitive data
- Acoustic emanations, acoustic signals that can leak plaintext data

Exploiting such emanations has long been of great interest to organisations concerned with intelligence gathering, and as such not much information is publicly available about recent and current implementations of related methods and technologies in those circles. However, in recent years there have been an increasing number of papers published in open academic literature investigating potential methods of attack and defence, in addition to open source work on both plaintext data processing and emanations published by private organisations and individuals [72, 86].

As stated before, the purpose of this project is to highlight the potential of these attacks by reporting on the advances in academic research on compromising emanations, especially focusing on the surveillance applications of such attacks. Their limitations and practicality will be assessed by looking at ease of implementation, what type of equipment would be needed and whether or not access to such equipment would be limited by cost and other factors. Current developments of technology likely to assist attack and defence scenarios, and possible future developments will also be examined and suggested.

Chapter 2

Electromagnetic Emanations

This chapter explores the development of published electromagnetic eavesdropping techniques. A short introduction covering their history places these attacks in context with their status as powerful tools for surveillance. The use of electromagnetic emanations for the reconstruction of computer display contents is covered, from the very first openly published attack to the latest methods against modern display technology and various types of keyboards, with suggested software and hardware countermeasures also being looked at. Finally, all this information is incorporated into an examination of the application of these methods and technologies for the purposes of both surveillance implementations and defensive measures.

2.1 Introduction

The history of the exploitation of electromagnetic (EM) emanations for information leakage stretches back to the forties; during World War II, Bell engineers noticed that compromising RF emanations from teletypewriters used for secure communications, using what was essentially a one-time pad, allowed the recovery of plaintext [66]. The subsequent drive to protect sensitive equipment from emanation based attack was codenamed "TEMPEST". Strictly speaking, the term TEMPEST now refers to a set of EM emanation related standards, aimed at defining methods and limits for the protection of electronic equipment and reduction of compromising emanations by US civilian and military organisations. Later on, during the Cold War, efforts to secure US military installations in Guam against compromising emanations lead to the implementation of TEMPEST standard protection measures, with the help of NSA COMSEC (Communications security) engineers, such as filters and improving equipment grounding [16]. Although some limited information is known about the TEMPEST standards, such as name and area covered, the majority of it is classified.

2.2 The development of emanation based attacks

With the growth of the security industry and open security research, TEMPEST rated equipment is not only available to private organisations and individuals, but actual studies of EM emanations have been undertaken in a great deal of depth in open, peer reviewed literature, taking the form of both new attacks and countermeasures. In fact, these studies, originally undertaken as an emulation of military and intelligence capabilities, can be considered as the progenitors of all open research on compromising emanations. To demonstrate their relevance, and the threat they posed to security, the attacks focussed on what was the latest technology at the time, personal computers, with computer displays and serial connections first being targeted.

2.2.1 Emanations from CRT displays

The first published paper revealing the potential for data recovery from EM emanations was published in 1985. The paper dealt with emanations from computer monitors, then based on cathode ray tube (CRT) technology [35].

In a CRT monitor the video signal is amplified from a low voltage for circuit level functions to several hundred volts when put through the CRT [35]. This results in the video signal being the major component of the emanation field produced by CRT monitors. This signal is similar to a television signal, and the success of this attack relies on the ability to reconstruct the monitor video signal using a television receiver. Synchronisation frequencies, horizontal and vertical, required for a stable image are part of the emanated video signal, but not recognised by the receiver. However, they can be generated separately by the attacker, and used to correct the video signal at the television receiver. One method of doing this was shown in the experiment, manually finding them by tuning an oscillator controlling horizontal frequencies in the range 15-20 kHz. Dividing the correct horizontal frequency by the number of display lines of the CRT gives the required vertical synchronisation frequency and this action can be automated by a digital circuit, so the attacker only needs to know the number of display lines and to manually adjust one oscillator in order for the television receiver to correctly output a stable recovered image. It was also suggested that recovering the synchronisation frequencies could be done by extracting the horizontal frequency from the video signal using a band pass filter and running this noisy sine wave signal through a pulse circuit. The frequency is then divided by the number of screen lines as before, and both synchronisation frequencies can be recovered.

The monitors tested did not produce emanations greater than IEC CISPR limits for data processing and office equipment in force at the time [35]. However, it was possible to reproduce a clear picture with a receiver 50 metres away from a monitor in plastic casing, and at about 10 metres for one in a metal casing. Effective distances could be improved on by using a directional antenna and amplifying any received signal, and it was estimated that signals may be recoverable from up to a kilometre away for monitors in plastic cases, and 200 metres away for those in metal cases. To replicate a practical attack scenario, the equipment was placed into a car which was parked outside a building where it was known a certain word processor

was being used. Static photographs of the recovered video signal were shown to clearly contain the same data displayed on the target monitor. The method was even demonstrated on BBC's "Tomorrow's World", a popular science and technology television show.

Countermeasures suggested by the author ranged from modifications to the design of CRT monitors to reduce the strength of emanations produced, to implementing shielding in the form of Faraday cages for monitors, which cannot cover the screen, and therefore only offer a partial solution [35]. Another possibility is the implementation of changing the display line order, so that the monitor can still display the correct image, but the television receiver cannot. Changing the display lines randomly over time can have the effect of the attacker not being able to reconstruct the image even with the correct synchronisation frequencies, and if they correctly guess line order, not having a stable image for any useful length of time. Interference purposely added by a monitor is not viable countermeasure, as monitor manufactures have to comply with standards to reduce emanation interference already. Neither is placing the monitor in a room with a large number of other monitors, as successful reconstruction of an image is still possible even across monitors of the same make and model due to distinct resonance frequencies.

2.2.2 Emanations from serial cables

Following on from the work of van Eck, it was found that RS-232 serial data cables are also capable of producing compromising EM emanations [78]. The emanations are high frequency signals corresponding to binary data rise and falls that occur as magnetic dipoles present due to capacitance around the cable when one of the connected devices is not grounded. It is worth mentioning that they are not present if both devices to which the cable is connected are grounded. Experiments conducted against an unshielded RS-232 cable found that bit data could be recovered from 7 metres away with a small short wave radio tuned to 16 MHz. Distances in other environments, and with other equipment, revealed that an attacker could expect to receive data signals in both the AM and FM band at about 6 to 9 metres, and even some shielded cables were still vulnerable. The author noted that this was a good enough attack to allow recovery of signals from a PC-modem cable through an adjoining wall in a semi-detached house. The low cost and small size of the necessary equipment also made this attack an potentially powerful method for eavesdropping in a domestic scenario, against a PC user in an adjacent house or apartment, for example.

These two papers caused great excitement in both academic and public circles. With a greater awareness of EM emanations being a security issue, many companies began selling van Eck type emanation receivers. However, greater academic discussion on further possible implementations of these types of attacks did not begin until later in the decade.

2.2.3 Emanations as a covert channel and defences against display eavesdropping

The paper that would further ignite academic interest in the security implications of EM emanations was presented in 1998 by Kuhn and Anderson [55]. Their work demonstrated practical

variations on the already known attacks, and also new methods of preventing data leakage, not by hardware modification but by software implementation. They first demonstrated how emanations could be used as a form of covert channel. By installing software designed to produce radio emanations from the CRT display on a target PC sensitive data can be broadcast to a short wave radio owned by the attacker, and a recording made for later analysis. The emanations are produced by using the monitor to display specific images that cause a desired CRT electron beam current. The emanations from this current can be recognised as a radio signal, and interpreted by an AM receiver as an audio tone. The images are a result of grayscale pixels displayed in an order determined by timing calculations, which are in turn determined by the properties of the monitor itself, i.e. horizontal and vertical frequencies, pixel clock and screen area desired for the image. The radio signal can be picked up with a simple radio and antenna from a close distance, which can be improved with better equipment, lower interference, and keeping the antenna close to the power supply cables that feed the monitor, which act as an antenna for the emanations. A data rate of 50 bits per second was achieved using this method.

The next experiment tested another method of transmitting recovered data covertly, this time in a way that would be very hard for a user to detect, even if transmission was occurring right in front of them [55]. Emanations were recovered with a commercially sold, TEMPEST influenced, EM emanation receiver, similar in design to the van Eck receiver originally proposed in 1985, against a modern CRT monitor. Covert data transmission was achieved by applying an imaging technique known as “dithering”, where high frequency patterns from a limited colour palette can be used to create the illusion of one solid colour not from the palette. This effect works fine on the human eye, but when eavesdropped via the receiver dithered areas of the display show up strongly due to the much greater sensitivity of the receiver to high frequency signals. To demonstrate the power of this method, it was shown how one word “OXFORD”, clearly visible to anyone viewing the CRT monitor as a solid colour, could be eavesdropped as “CAMBRIDGE”, this word being present in the original image as dithered grayscale characters surrounding the first word, visually indistinguishable to the user, but showing up strongly on the eavesdropped signal.

The last experiment demonstrated the effectiveness of a novel countermeasure against the recovery of text, named “Tempest Fonts” [55]. By taking advantage of the fact that an EM receiver can only display signals of a high enough frequency, it is possible to use a specially designed font that is clear to the user, but does not create a recognisable signal and image of text on an attackers receiver. This idea was based on the observation that only the top 30% of the horizontal frequency of an emanation signal is actually recognised and displayed by the receiver. By running the pixel field of some displayed text through a low pass filter to remove the top 30% of the Fourier transform of the horizontal frequencies, it was found that the received signal at the EM receiver did not display any recognisable text, even when placed right next to the monitor. The difference in appearance of the altered font on screen is barely noticeable to a user, only appearing slightly fuzzy when scrutinising pixel displays. The authors noted that the exact proportion to be filtered would vary according to equipment used, but that 30% was a good yardstick for further measurements. The ease of implementation of this defensive measure, its low cost, and its effectiveness made it an ideal addition to already existing hardware countermeasures.

Repeating the above attacks using a TEMPEST rated receiver improved the power of the attack, and also allowed the implementation of automatic character recognition from averaged images [52]. Recognition rates for this feature could achieve around 66% from images averaged from 16 frames, but were highly sensitive to low signal-to-noise ratios. Implementing the "Tempest Fonts" also prevented useful levels of character recognition.

The successful demonstration of eavesdropping attacks against standard consumer grade CRT monitors underlined just how vulnerable computer systems could be to attacks exploiting EM emanations, even without physical and network access. Serial cables and CRT displays, however, were beginning to be phased out by manufactures, in favour of better, more advanced alternatives, leading to the development of new attacks on more modern equipment.

2.3 Emanations from more modern equipment

The continued relevance of EM attacks has been demonstrated in the successful application to emanation eavesdropping techniques to newer display technologies. The successful implementation of these attacks against LCD monitors has also led to the continued development and deployment of the software countermeasure "Tempest fonts", originally developed for CRT monitors.

2.3.1 Emanations from LCD displays

By 2005, CRT monitors were on their way to being replaced by liquid crystal display (LCD) monitors for personal computing use. Because LCD monitors do not use the same EM emanating components that CRTs do, such as magnetic deflection coils, and they operate at a lower voltage, it may seem that they would be less vulnerable to the EM eavesdropping techniques discussed above. However, it was demonstrated that some LCD monitors can be just as susceptible to EM eavesdropping techniques as their predecessors were [53]. Experiments against two LCD displays, a laptop screen and a desktop monitor, were undertaken using an EM receiver, designed for testing emanations to the confidential TEMPEST NACSIM standards, capable of receiving frequencies of up to 1 GHz with intermediate frequency filter bandwidths of 50 Hz to 200 MHz. This is important because the inverse of the filter bandwidth determines the shortest impulse that can be used as receiver output. To successfully eavesdrop an LCD display it is necessary for this value to approximate the pixel clock frequency of the display, with the pixel clock being necessary to correctly sequence pixel values on the display. The filter bandwidths of AM receivers and TV tuners are usually only capable of implementing bandwidths of up to 8 MHz, and as such, they would not be suitable for implementation as EM receivers for attacks on LCD displays. To fully reconstruct the image, the receiver signal was output to a digital storage oscilloscope with special software being used to convert the signals to images. As well as testing for information leakage, methods for mitigating the vulnerability were also suggested and analysed.

The first experiment against the laptop display resulted in a readable image being produced from measurements taken from 3 metres away, the highest quality images being found at a frequency

of 350 MHz with an intermediate bandwidth frequency of 50 MHz [53]. Highly legible images were produced by averaging multiple images together to reduce noise, however, single images with full background noise still contained readable text. By using an improvised EM probe, made out of a common coaxial cable, it was possible to ascertain that the emanations were not coming from the laptop display itself, but instead coming from the cables connecting the laptop motherboard to the display. These cables form part of the Flat Panel Display Link (FPD-Link) interface for connecting video controllers to displays, and is used in most laptops. The FPD-Link is an implementation of a low-voltage difference signalling (LVDS) system. LVDS systems use voltage differences across twisted pairs to transmit information. In a laptop FPD-Link, twisted pairs transmit pixel colour information, horizontal and vertical synchronisation and a control signal. One twisted pair is used to transmit the pixel clock frequency with the bit rate for such a connection being seven times the clock frequency.

The use of twisted pairs in LVDS is an attempt to minimise interference from data transmission [53]. However, emanations are still recoverable from such cables, as the results of the experiment show. The measured field strength of the signal received from the laptop from three metres away, while displaying a boot screen with default colours, was thought to be of sufficient strength to allow eavesdropping from several rooms away. To demonstrate this, the next experiment using the laptop was conducted from 10 metres away, through three plaster board walls. Legible screen images were attained through the application of the image averaging technique mentioned earlier. This was applied automatically by calculating the necessary frame rate through cross-correlating the first and last image recorded. It was noted however that in practice this method might not be implementable, as frame rate selection can be incorrect where interference is present. In this case it is possible to manually select the correct frame rate. The lower value of the measured field strength at this distance was found to correlate with the increase of distance and obstacles presented by the plaster board walls.

To determine factors influencing legibility of screen text in reconstructed images, and the possibility of a countermeasure for preventing successful eavesdropping, various styles of text were tested. These can be seen in the test image in Figure 1. Essentially, the different types of text were;

- Black text on white background
- Maximum contrast in grayscale and colour
- Minimum contrast coloured text and backgrounds
- Three lines with random bits added to pixel values, except text

It was found that all text, apart from the lines which had random bits in them, were at least somewhat readable across two frequencies; 350 MHz and 285 MHz [53]. Figure 1 demonstrates exactly how differences in the style of text result in changes in legibility at the receiver. As can be seen, the use of random bit values was the most effective method of introducing illegibility into the received signal. However, it was noted that to be effective, new random bit values had to be selected per character each time. It was determined that reuse of the same values

Figure 1: Comparison between test graphics and received image from a laptop LCD display.

This figure demonstrates the effectiveness of different text formats, the last three lines of both images using random pixel values as a potential defensive measure against information leakage. Images taken from [53].

per character could in fact aid an attacker using an automated character recognition scheme, since the addition of repeated random bits would increase the uniqueness and identifiability of a particular character with regards to other characters undergoing the same treatment.

Next, the author tested the eavesdropping potential of a desktop LCD screen [53]. The display used in the experiment uses a standard for relaying information from the video controllers to the screen called the Digital Visual Interface (DVI), and there are two others in use by other models of LCD display. These standards all use an underlying technology called Transition Minimised Differential Signalling (TMDS). Like LVDS implemented in FPD-Link, TMDS uses differential signalling across twisted pairs to minimise EM interference, transmitting information as encoded bits. It is an improvement over LVDS, however, in that the transmission algorithm is also designed to minimise interference by implementing DC-balancing. Using similar test images as shown in Figure 1 above, several styles of text were eavesdropped to determine legibility and for comparing against the results from the laptop display. Text was generally much less readable than that which was recovered from laptop emanations, minimum contrast text being the most recognisable. Randomisation of bit values was again the most effective method of preventing successful eavesdropping, with TMDS encoding lending itself to this method by requiring a lower number of random bits to achieve the desired effects, making the displayed text easier on the eyes of the user, while still preventing eavesdropping.

All in all, it was demonstrated that LCD displays were capable of emitting compromising emanations. The legibility of eavesdropped text could both be improved or reduced by careful selection of the method of displaying the text, an attacker could force the use of highly contrasting text for increasing legibility for example, whereas someone trying to protect a system could choose to add random bits to each character display to prevent the readability of text in any emanations.

2.3.2 Improving on the "Tempest fonts"

Further work on inducing illegibility of eavesdropped screen images has resulted in improvements on the original "Tempest Fonts" [82]. By running the pixel field through a Gaussian filter, in addition to the Fourier transform, it is possible to decrease legibility compared to text that has solely undergone Fourier transform filtering. This benefit, however, comes at the price of further decreasing the legibility of text on the display itself.

More recently, another novel method for preventing successful image reconstruction from emanations, the source being analogue video cables connected to both CRT and LCD monitors, has been proposed [90]. By taking advantage of how the human visual system processes information, it is possible to alter the displayed image in a manner imperceivable to the user, but which renders the eavesdropped image illegible. The system works by adding random pixel values to the intended display, and at the same time removing the same corresponding pixel values from the original image. Displaying these images in succession results in a user effectively "seeing" the original image, while eavesdropped images appear heavily distorted. While some vestiges of the original image may still remain, the manner in which this countermeasure can be implemented, in this case with a peripheral FPGA (Field-programmable gate array), means that it can be combined with the aforementioned "Tempest Fonts", to potentially eliminate any trace of text recovered by an attacker.

These measures offer a good level of protection against EM emanation surveillance of data output, but keyboards, the ubiquitous method of plaintext data input, are also vulnerable to attack.

2.4 Emanations from keyboards

Keyboards have also been found to be susceptible to EM eavesdropping, allowing remote keystroke detection [88]. This paper demonstrated four methods of implementing such an attack, against a wide variety of keyboards using PS/2, USB, laptop and wireless connections, in a range of environments. The keyboards were attached to laptops running on battery, so that emanations were not introduced into the shared ground, provided by the power cable connecting to the power line, meaning that all the emanations tested in the paper, unless mentioned otherwise, are purely being picked up from the keyboard itself. One aspect of note is that the authors demonstrated that the use of commercially available EM emanation receivers was not necessary and that using open source software radio technology to perform the task of an EM receiver

was possible. However, they found better results could be achieved using another method. By taking the raw signal capture from an antenna and applying a short time Fourier transform, displaying both frequency and amplitude over time, they could automatically identify emanations from voltage peaks, which act as "triggers", using radio analysis software. However, this can not be done in real-time. The signal from the antenna first has to pass through an analogue to digital converter. This technique results in a very high sample rate, however, and there was no sufficiently fast way of transferring the resulting digital signal to a PC. The authors instead fed the signal through a digital oscilloscope that was able to detect potential emanation peaks and pass just these signals to a PC for further analysis. This allowed the capture of all EM emanations of any frequency up to 2.5 GHz without having to manually select a frequency for demodulation by a receiver.

The authors then used this collection method to implement the four techniques in experiments against various keyboards [88]. The first technique, called "Falling Edge Transition Technique", tested the recovery of partial keyboard scan codes, which encode which key was pressed, from the falling edge of the signal generated by each keystroke on a PS/2 keyboard from five metres away, in a low interference environment. The clock and data signal of the falling edge, expressed as peaks, can be used to identify likely keys by grouping signal traces with partial scan codes. This allows a partial recovery of keystrokes, and can be used to improve a dictionary or brute force attack.

Improving on this method, the second technique, called "Generalised Transition Technique", tested used a bandpass filter to improve peak detection of the raw signal from PS/2 keyboards [88]. This allowed recovery of the rising edge of the each signal, these being present between peaks which result from both the clock and data signal simultaneously. With knowledge of the rising edges, in addition to the falling edges, of the signal, it is possible to determine the full value of the scan code, and thus determine exactly which key the emanations are from.

The third technique, "Modulation Technique", utilised additional information from harmonic interference, identified as most likely coming from the keyboard's microprocessor clock or crosstalk between the data and clock signals of PS/2 keyboards [88]. When the harmonic signals were compared to the data and clock signals it was found that the harmonic signals contained all the information provided by the others. This allows recovery of the scan code from harmonic signals, and also their use in conjunction with data and clock signals to improve the signal-to-noise ratio of readings, resulting in better measurement accuracy. Also, because the harmonic signals are frequency and amplitude modulated, they are more resistant to other noise interference and physical obstacles, so the effective range for a practical eavesdropping attack is increased. However, to capture a signal, a demodulation step has to be added to the previous methods using an oscilloscope.

The final method, called "Matrix Scan Technique", is applicable to all types of keyboard, not just those that use a PS/2 connection [88]. This technique uses emanations from the keyboard matrix to determine keystrokes from a group of keys linked to the same matrix column. Matrix emanations are different from the other emanations exploited before. The design of a keyboard matrix means that each key is at an intersection of columns and rows. The keyboard controller sends pulses down each column in turn, with a detector for each row. Once a key is pressed down

and the pulse is sent down the correct column, the corresponding row is detected. Used like coordinates, the scan code for the letter can be found and transmitted to the PC. It is worth mentioning that the specific details of which keys belong to which columns will vary across makes and models. The attack works by eavesdropping emanations caused by the column pulses, which produce a continuous signal, punctuated by voltage peaks when a key is pressed. The total number of peaks detected indicates a particular column has been activated. By grouping keys into their respective columns, it is possible to determine by the number of peaks which column the keystroke belongs to. Practically, this allows partial keystroke recognition much like the first technique demonstrated.

After demonstrating that keystrokes were recoverable using EM emanations, the issue of how to identify emanations from one specific keyboard among others was also addressed [88]. The first three techniques make use of clock signal frequency, which tends to vary uniquely across makes and models. By checking the frequencies of received clock signals against the time lapse between them, it is possible to differentiate between two keyboards. For an attacker using the "Modulation Technique" differences in clock frequency signals from the microprocessor can also be used as an identifying factor. When using the "Matrix Scan Technique" it is possible to use the emanations to establish continuous use of a keyboard, due to the regularity of the column pulses. Peak duration also varies across keyboard models, so this may also be an identifying feature. Physical characteristics of each keyboard may also be used; PS/2 cable acts as an antenna, different lengths affecting signal strength, laptop cables are shorter, resulting in a higher frequency band and noisier signal, the short length of keyboard matrix wires also result in a higher frequency.

To demonstrate the effectiveness of the above techniques against keyboards in a real world scenario, experiments were undertaken in a variety of environments, on several different types of keyboard [88]. A successful attack was classified as one that achieved a 95% recognition rate of keystrokes from a 500 word sample. The results showed that the maximum range of a successful attack depended on the keyboard that was being tested. In an office environment, with line-of-sight access to the target keyboard, the maximum range of the "Falling Edge Transition Technique" was found to be around 3.5 metres for the least vulnerable keyboard, and 10 metres for the most vulnerable keyboard. The ranges for the "Generalised Transition Technique" were 3 metres for the least vulnerable keyboard and around 7.5 metres for the most. For the "Modulation Technique" the ranges were 5 metres and around 9.5 metres, and for the "Matrix Scan Technique" they were around 1.5 metres and 3 metres. The authors found that the results for attacks from an adjacent office, separated by a wood and plaster wall, were similar, but ranges were decreased due to the lower signal-to-noise ratio caused by the physical obstruction. One interesting result was produced when attacks were attempted from within a building, where effective ranges of up to 20 metres were achieved. The authors reasoned that the building's main electric line, being the shared ground, acted as an antenna, making the real range of the attack the distance from the keyboard to the electric line, which needed to be less than one metre for a successful keystroke recovery, and the distance of the electric line to the receiver. They did find, however that the "Matrix Scan Technique" was heavily affected by noise and interference on the line, and they were unable to achieve stable measurements using the technique. They also found that the main water pipe was a viable alternative, since it was less affected by interference, but they did not mention whether the last technique worked using the

water pipe or not. Using a shared ground in the office environment itself did not work however, due to too much noise. It did work, however, with a probe attached to both the keyboard and the ground, with a receiver successfully picking up emanations from the ground. They also noted that the receiver could be directly hooked up to the ground, but that this type of attack ceases to function as emanations eavesdropping, being more akin to power analysis.

Countermeasures suggested by the authors included implementing EM shielding of keyboards and rooms where secure computing needs to be conducted [88]. These are very expensive measures however, and may be beyond the implementability of many individuals and organisations. One way to prevent the attacker from recovering meaningful data from emanations is to encrypt the serial data link from the keyboard to the computer. However, this does not provide protection for data eavesdropped using the "Matrix Scan Technique". The authors suggested that an extension of the countermeasure suggested by van Eck, changing display line order to protect display emanations [35], could be implemented; by randomly running through the column order of the matrix scan, and randomly adding delay to the microprocessor subroutines it may be possible to obfuscate keystrokes so emanations yield less information to an attacker [88]. One other method, detailed in a patent [68], involves using high frequency filtering of matrix signals before processing by the keyboard's microprocessor to reduce EM emanations from the matrix itself.

Several possible extensions to this attack were also mentioned; improving keystroke detection through sophisticated filter techniques or by acoustic detection and improving the achieved computation time of 2 seconds per keystroke through the use of dedicated hardware, such as FPGAs [88]. It was also mentioned that by using software radio technology they could make the attack portable, requiring only a laptop and a clandestine antenna hidden underneath the attacker's clothes. This came at a price however, as the maximum range of such a set-up was found to be two metres, though distance may not be an issue for an attacker able to physically approach the vicinity of the target system.

Indeed, the highly specific nature of EM emanation based attacks lends itself to the nature of targeted surveillance. The knowledge and equipment needed for the successful implementation of an eavesdropping attack was previously limited to individuals and organisations with the necessary authorisation and funding, but with greater awareness and technological advances in open forums this is no longer the case.

2.5 EM Emanations and surveillance

2.5.1 Methods of attack

While some attacks outlined above have used TEMPEST style and TEMPEST rated receivers in order to demonstrate eavesdropping capabilities [52, 53, 55], others have shown how commercially available, and also open source, spectrum analysers and radio receivers can also be used to successfully recover data from compromising emanations [82, 88, 90], equipment that is freely available to any interested party. The mobility of such attacks has also increased; whereas

the computational analysis of recovered emanations required static computer equipment, and receivers and antenna equipment could only be fitted into cars and vans requiring a team of attackers, it is now possible to perform automated emanation analysis on laptops, easily carried and deployed effectively by one person. Prior knowledge of a targets system is still essential for an attack to work; make and model of monitor, or keyboard, and type of connection used, for instance, but this information would usually be gathered or inferred during a surveillance operation anyway.

The capabilities of well funded organisations are likely to greatly surpass what has been proposed in open literature. Ongoing DARPA (US agency; Defence Advanced Research Projects Agency) programs such as ChaSER [24], a project to develop ELINT (Electronics Intelligence) and SIGINT (Signals Intelligence) receivers for use on UAVs (Unmanned Aerial Vehicles) and unattended ground sensors to increase remote surveillance capabilities, and RADER [29], which will vastly improve the capabilities of analogue to digital converters, equipment that is vital in exploiting EM emanations, are at the cutting edge of electronic surveillance technology, the implementation of which will no doubt spur on the development of countermeasures to defend against attack and prevent compromising emanations.

2.5.2 Methods of defence

Defending against emanation based attacks can involve implementation of both physical and computer security measures. Perhaps the most obvious countermeasure is the use of a Faraday cage, with protected power line and data cables, to contain sensitive computing resources. While providing a high degree of security, this solution is inflexible and would not allow mobile access to sensitive information. Building on the concept of shielding, there are many private companies offering what they claim to be TEMPEST rated computing equipment, including monitors, desktop computers, laptops and phones [22, 77, 79]. There is undoubtedly a high cost associated with these devices; the fiscal overhead costs incurred by the manufacturers in design, materials and standards compliance, is something which is ultimately passed on to the buyer.

These countermeasures are easily implementable by any organisation with the resources to do so, but for smaller organisations and individuals they are not practical. There have been calls for manufacturers of electronic equipment to incorporate design techniques and technology into their products for the purpose of reducing the security risk presented by compromising emanations, but it has been shown that for this to be achieved, manufacturers would have to vastly improve on the measures already in place, which are designed to comply with current required EM standards [54]. The manufacturing industry's motivation for reducing emanations from consumer grade equipment is purely economical, they would not be able to sell equipment unless it complied with standards designed to ensure that no excess interference is caused by the product, and as such, manufacturers cannot be expected to produce secure equipment for the general public without incentive. Implementation of the previously mentioned "Tempest Fonts" is offered by security companies as part of commercial software security suites, PGP corporation first included a "Secure Viewer" in PGP 6.0.2 [69], and it is still available in current PGP products. However, this countermeasure only protects what is covered by the secure

viewer, it does nothing to prevent or reduce the existence of EM emanations, such as those from keyboards. That being said, there is nothing to stop any user attempting to implement home-made solutions, previous demonstrations on the effect of aluminium foil on RFID jamming of passports [60] may be grounds for a small study to test the effectiveness of ad hoc emanation blocking measures directed at vulnerable equipment, for instance home-made enclosures for monitors and keyboards of shielding for the analogue video cables of LCD monitors.

Chapter 3

Optical Emanations

In this chapter, attacks exploiting optical emanations are examined. Basic forms of optical attacks are first looked at, with real world instances highlighted. Published attacks involving novel methods of reconstructing computer display content from optical phenomena, such as diffuse emanations and reflections, and recovering data from direct imaging of sensitive information, flickering LED status lights, keyboard input and physical key shapes, are then reviewed in depth. Finally, the continued development of advanced optical equipment is discussed in relation to their potential use in optical attacks, and possible defence methods against this type of surveillance are explored.

3.1 Introduction

Traditional surveillance tools bring to mind binoculars and telescopes, tools with which to spy on someone, or something, from a distance, to gain a visual understanding of what is going on. Optical attacks are an extension of these methods, allowing the attacker to gain knowledge of sensitive, perhaps cryptographically protected, information from data entry or displays. The simplest of these methods is shoulder surfing, where the attacker compromises the security of sensitive information by looking at it surreptitiously from a nearby location such as over the shoulder of the legitimate user. This has long been known to be a security issue for all cryptographic implementations requiring user input for authentication, such as password or PIN entry [3]. This basic method can be improved by the use of traditional surveillance tools such as those mentioned above, or more modern equivalents, such as concealed video recording or relaying devices, something which has increasingly been implemented by criminal gangs against ATM PIN entry interactions in recent years [9]. In addition to data entry, the displaying of protected data in human readable format is another viable avenue for attack.

With regards to the term "optical emanation" one may consider that all light emitted and reflected from physical objects carries information usable by the human visual system, so indeed any image is an optical emanation of sorts, and surveillance attacks making use of such emanations

are directed towards obtaining a useful image or images of sensitive data, such as a displayed document or PIN entry. Recently more novel methods of retrieving information from targets via optical techniques have appeared in the academic literature, concerning the retrieval of both sensitive displayed data, and sensitive user input.

3.2 Novel optical attacks

Unlike the development of EM eavesdropping attacks, optical attacks have not generally followed on from each other. Though they share the use of similar imaging tools; photosensors, cameras and telescopic equipment, each of the attacks reviewed here takes a different approach to reconstructing an image of sensitive data displayed on a monitor or input on a keyboard; by making use of light and image reflections, decoding data from status LEDs or direct imaging of sensitive data.

3.2.1 Reconstructing an image from diffuse optical emanations

One of the first novel optical attacks exploited emanations of displayed data from CRT monitors, emanations which can be analysed without line-of-sight access to the display [51]. CRT monitors, which today are disappearing from widespread use, employ a method known as raster scanning to reconstruct images from a cathode ray tube onto a monitor line by line, the timing of which is controlled by the pixel clock frequency, which is the multiplicative inverse of time taken for the electron beam to travel from the centre of one pixel to the centre of the pixel to the right. This clock value also determines the horizontal and vertical deflection frequency, which in turn determines the rate at which the scan lines are drawn and frames are built on the screen, respectively. Knowing these deflection frequencies allows an attacker to construct a stable image from these emanations. The emanations themselves are the diffuse light reflections from the immediate environment around a CRT monitor, such as from a wall or through frosted glass.

The intensity of the light emitted by all pixels on a CRT monitor is the equivalent of the video signal convolved with the impulse response of the monitor phosphors [51]. The attack method utilises measures of the light emanation intensity to attempt to recover the video signal. Knowledge of the phosphor decay times of the phosphors used in the monitor is important though, as it can be used to improve image quality extracted from the recovered signals. As such the decays were measured for the monitor used in the attack. However, it may be possible for an attacker to use manufacturers references to pre-emptively determine decay times, depending on availability.

The attack itself produced some very interesting results. From a distance of 1.5 metres, and in a dark room, a photomultiplier photosensor module was placed facing a wall. In front of it, facing away from the photosensor, was a CRT monitor displaying a test image of white text of varying sizes on black background, with test letters in red, green, cyan, magenta and yellow, at VESA standard 640x480 resolution at 85 MHz [51]. Using the photocurrent detected from the

photosensor and decay times to extract the video signal, a grey scale image of the test image was reproduced. The largest sized text was readable, but each white pixel appeared smeared across the image due to the decay picked up from the raw signal. Running the signal through a Butterworth high-pass filter cutting off frequencies lower than 4 MHz greatly improved the quality of the image, with the most of the smaller text readable, and the coloured text also readable. An even better image was produced by deconvolution, this was achieved by applying Fourier transforms to both the average luminosity signal and phosphor impulse response per frame, dividing these results and then applying an inverse Fourier transform. This process allowed for even the smallest font to become readable, with each character appearing much sharper compared to the previous results.

A theoretical analysis of the threat potential of this kind of attack was also undertaken [51]. For an attacker with line-of-sight capabilities it was estimated that a signal could be recovered, from a CRT in a well-lit office environment, from up to 80 metres using a simple telescope. Without line-of-sight capabilities, an attacker could theoretically detect a signal, from a reflecting wall in "late twilight" light conditions, from a distance of up to 50 metres using the same type of telescope as above.

The author made several suggestions on how to improve the strength of this attack through the use of other equipment, such as the use of more powerful telescopic optics, like a telescope with a zoom function, to allow the greater isolation of light coming from the target which would increase photon capture and improve the signal-to-noise ratio [51]. Also suggested was a narrowband photosensor to improve raw signal quality and improved filtering techniques such as analogue preprocessing to better approximate a better deconvolution filter than the one used in the experiment.

Possible countermeasures suggested by the author included;

- Keeping CRT monitors that display sensitive information away from windows and other places that allow direct and indirect observation
- Using background lighting that greatly adds noise in the spectrum of CRT phosphors such as incandescent lighting
- Lowering CRT luminescence during lower ambient lighting, which the author notes would also save power and users' eye strain

Finally, the author noted that the, at the time, increasing use of LCD monitors would reduce the need for security considerations for CRT use, as LCD monitor pixels are much slower than CRT phosphors and LCD monitors update all pixels in one line at a time, instead of a sequential CRT raster scan line, indicating that LCD monitors are at a much lower risk from the types of optical attacks presented in the paper [51].

3.2.2 Recovering data from status indicator LEDs

Another novel attack makes use of a different, just as ubiquitous, feature of modern computing for information leakage and transmission applications; LED status indicator lights [61]. Based

on observations of correlation between changing RS-232 serial communication signal voltage levels and modulated optical emanations from an LED indicator, the authors investigated the possibility of developing practical attacks against a wide selection of computing and networking equipment. The correlation is due to LEDs usually being connected to the logic gates that make up a logic circuit, this results in a direct representation of the bit stream of the serial connection appearing as flickering of indicator lights. Due to the fast response properties of LEDs, the serial data signal can be reproduced as an optical emanation at around the same data rate as achieved by the serial connection. The authors noted that the high noise levels that can be encountered when attempting this attack can be mitigated by using a sensitive detector and telescopic optics, where the LED can be focussed on completely to reduce background light radiation as much as possible. The signal can also be fed through an off-the-shelf Universal Synchronous/Asynchronous Receiver/Transmitter which allows the data to be recovered from a noisy signal.

The authors designed a classification system for optical emanations in real world implementations. Classes I, II and III are, in ascending order of severity, measures of just how much information is leaked to an attacker from a indicator light [61]. Class I indicator lights merely expose the equivalent of one bit of information, such as a power-on light shows that a machine is on. Class II indicator lights give out more information, such as the amount of data, which opens up avenues for traffic analysis. Class III indicators generate emanations that are highly correlated with the data that is being transmitted or received, which may allow for the retrieval of data by an attacker. The authors then conducting an experiment on 39 assorted pieces of equipment, with a total of 164 LED indicator lights, examining each light closely with a PiN photodiode detector and oscilloscope and comparing this to readouts from a breakout box hooked up to the serial connection, to see if any of the LEDs exhibited class III behaviour. They found that 14 items exhibited class III behaviour, mainly routers and modems; communications devices. One interesting observation was that there was no significant difference in the readings produced by smaller surface mounted LEDs compared to standard LEDs, with their brightness being comparable, which suggests that LED brightness is a large factor in signal strength.

An experiment was then conducted, designed to test for feasibility of a practical attack on class III LEDs by examining how distance from the LEDs affected signal measurement [61]. By taking measurements under variable light conditions, from incrementing distances, of LEDs on equipment transmitting and receiving data at variable data rates the authors found that correlation between data and modulated optical signal remained high to about 10 metres, after which the signal deteriorated heavily, where at 30 metres there was no significant difference in detection between a real signal and a random one. It was also found that faster bit rates did not make a difference to correlation, with some devices producing detectable correlations at data rates exceeding their stated maximum.

Special mention was made of one device, the InfoLock 2811-11, a standalone DES encrypting unit for use on financial wire transfer and ATM networks [61]. A related model, the 2811-13, had received Federal Standard 1027 endorsement by the NSA and was rated for purchase and implementation by federal departments and agencies until mid-1997 [65]. They found that class III LEDs were placed on the terminal side, coloured red, which took incoming data, later encrypting it, and sending it out the communications side, coloured black, corresponding to the

well known red/black concept of plaintext/ciphertext separation [61]. This was identified as a serious design flaw, using an optical attack it would be possible to reconstruct all the plaintext data from the LED as it came into the device before it was encrypted.

Overall, the authors had succeeded in demonstrating that optical emanations produced by LEDs were capable of leaking information. They concluded that although noise had been a large limiting factor in their experiments digital signal processing technologies can help reduce this problem, using a low pass filter to remove sub 120-Hz noise from the optical signal [61]. They also believed that the results from their long range experiments could be improved upon with improvements to both the optical and detector-amplifying components of the attack equipment; with better optical equipment they believed that they could recover optical signals from a few hundred meters away while using the same detector. A practical upper boundary on data rates by optical emanations was estimated at 10 Mbps (Megabits per second), but they thought that greater data rates may be feasible.

Suggested countermeasures against this sort of attack ranged from the pragmatic; such as moving vulnerable equipment away from risky line-of-sight areas or covering LEDs with black tape, to more technical and wide ranging solutions; using a technique called pulse stretching to minimise information leaked to attacker [61]. This takes advantage of signal jitter intolerance by making the minimum LED "On" time 1.5 times the unit interval of the data rate of the serial connection, the most extreme variation of this countermeasure makes the "On" time equal to the value of the current data rate or the slowest available. These countermeasures will essentially make what would previously have been a class III LED transmit the same amount of information as a class II LED.

The authors also investigated the potential that optical emanations hold for use as a covert channel [61]. They focussed on utilising emanations from LEDs on a keyboard, its three standard LEDs; "Num Lock", "Caps Lock" and "Scroll Lock" being easy for an attacker to locate, to relay keystroke information from both a software keylogger and a modified keyboard, but in principle this channel could work with any device or component with LEDs that the attacker can gain control over. They first wrote a software implementation of their attack, which could take control of all three LEDs, with a single LED being able to relay 150 bps (bits per second) and all three, if modulated by the same signal vastly increasing brightness and therefore range, but if each transmitted its own bit in parallel the method could achieve around 450 bps of ASCII data. The authors were also able to make this channel function on a wide variety of operating systems and keyboards.

They then tried modifying the keyboard hardware directly, and found that they achieved better results. By connecting the Scroll Lock LED to the keyboard data line and inverting the data signal to the LED they succeeded in modulating the LED with the keyboard data stream [61]. The effect for the user is hardly noticeable; the LED flickers a little with keyboard activity but when idle it is off, Scroll Lock still functions as a key but the LED does not come on, and when the keyboard goes through the POST stage the LED comes on like the others on the keyboard. This method achieves data rates equal to that of the serial connection of the keyboard itself, and relays data in the form of keyboard scan codes, which will require decoding by the attacker but has the advantage of containing more information than plain ASCII, it also has the potential to report key timing statistics, since both key press and release data is transmitted.

Potential improvements to this attack were suggested by the authors, most intriguing being the addition of an infra red LED combined with the pre-existing one on the keyboard [61]. This would allow IR detection methods to be used, possibly reducing the effect of a noisy environment, depending on conditions. It would also be entirely undetectable to the naked eye. The authors noted that while on the topic of physical modification on the keyboard even more effective measures could be put in place, even suggesting encrypting and compressing the serial connection data before transmission through the LED, improving efficiency and preventing a third party snooping on the established covert channel.

While both of the above attacks make use of photosensors to detect changes in emitted light, to help reconstruct an image thereafter, other methods making use of optical equipment to improve on the standard shoulder-surfing type attacks can be implemented with more widespread photographic and telescopic tools.

3.2.3 Image capture from reflections of a computer display

Greatly expanding on the most straightforward form of optical attack; imaging of a computer monitor directly, work has been done on recovering information from reflected images of the monitor found in various objects around the screen, including a human eye [6]. The attack makes use of low-cost consumer-grade digital SLR photographic and telescopic equipment; cheap refractor and Newtonian telescopes, easily available to any potential attacker, allowing the recovery of images from distances of up to 10 metres, while a much more expensive high quality Newtonian telescope can be used to capture readable images of large font and pictures from 40 metres away. It should be noted that this telescope, apart from being prohibitively expensive, is very bulky and heavy and it is unlikely that a set-up such as this could be implemented in a real practical attack.

Experiments using the low-cost equipment produced good results from objects with large, smooth surfaces. The best results were from teapots, from a distance of 10 metres the 18 pt font of the test image was readable in two of the three teapots tested, notably the two successful recoveries were performed on clear glass teapots filled with dark tea, perhaps providing a good contrast that increased image clarity [6]. Most impressively, from a distance of 5 metres, it was possible to recover 12 pt font from a reflection of an open word processor document. Other objects that were found to reflect usable images, with readable 18 pt font from 5 metres, included;

- Wine glasses
- Spectacles, even if they had anti-reflective coating
- A metal spoon

The authors found that images taken from outside when temperatures differed greatly were affected by air currents around poorly insulated windows which caused blurring in images taken using longer exposure times. They suggested that this could be circumvented by taking

multiple short exposure pictures and then combining the best images to produce a single sharp image, a technique used in astrophotography.

Smooth reflective surfaces were vital to recovering a good image, since it was found that the uneven surface of a plastic bottle meant that the image produced was somewhat distorted and some text was unreadable [6]. Also tested was the reflection of a paper A4 print on a teapot situated adjacent to it, from a distance of 5 metres. This produced an image which contained readable 10 pt font. The authors had some difficulty extracting a usable image from a human eye due to motion blur caused by eye movement, due to the reflection being quite dark long exposure times are needed, these factors combined place a limiting factor and made a practical attack using a corneal reflection impractical, as can be seen in the left image of Figure 2. They did note, however, that the cornea had excellent reflective properties, and found that "ideal" pictures taken from close by produced highly readable images limited only by the resolution of the imaging equipment being used.

After analysing the results of their experiments the authors found that limiting factors for this type of attack came in two main forms, the angular resolution of the optical device being used, that is to say the minimum aperture needed to capture the full resolution of the image, and the exposure time of the device, which adversely affected image sharpness when imaging a non-static object [6].

This attack was followed by an improvement published a year later [7]. The authors made reflections taken from a human eye their focus, due to the certainty of an eye being in the vicinity of the display of sensitive information. The equipment used in this experiment was of a lot higher standard than what was used in the experiments conducted a year previously, an astrophotography camera was used which cost 7.5 times as much as the DSLR, and a better quality telescope with a larger diameter, which was also smaller than the smallest Newtonian telescope used in the first low-cost experiments in the previous paper. Improved image quality in these new experiments was partly due to the higher quality of equipment used, but also due to the application of an image deconvolution algorithm, the Richardson-Lucy deconvolution, which was needed due to the very small depth-of-field given by the telescopes. This tends to produce blurry images if the correct focus is not found with a small margin of error.

The algorithm works by calculating the most likely pixel value from a known point spread function (PSF) and the value of the observed pixel, from the blurred image [7]. The PSF can be measured both before an attack, or during. Using a series of PSFs measured before the attack has the benefit of allowing for more accurate measurements and is able to produce very high quality images for static reflections, however, the method performs poorly when motion blur is present. The other method makes use of a reference point of light during image capture to extract the PSF. This provided a much better quality of image when used for image retrieval from a human eye, the authors also noted that reference lights for PSFs could come from a number of likely sources in a practical attack, such as status LEDs on electronic equipment, or a light source viewable in the background.

Using the above methods the authors were able to retrieve just about readable 18 pt text from an image reflected in a teapot from 30 metres away [7]. They were also able to retrieve an image

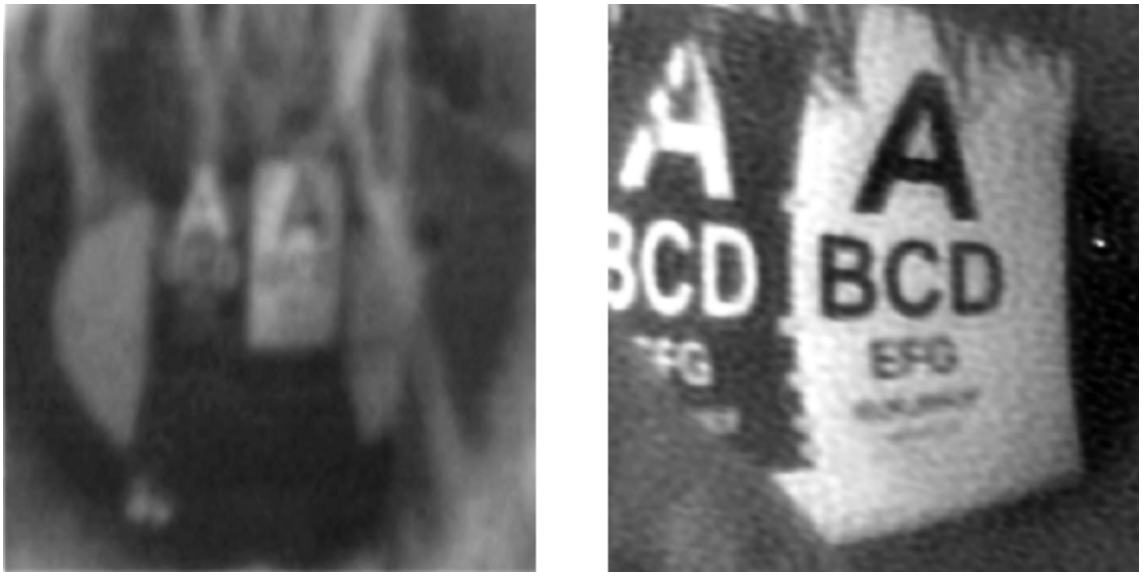

Figure 2: Comparison between two methods of optical reflections attack.

The older attack is on the left, taken from 3.5 metres [6] and the newer attack is on the right, taken from 10 metres [7].

from a human eye, taken at 10 metres. Figure 2 demonstrates just how much more powerful the new attack method is, achieving readable 36 pt font, a vast improvement over the previous version of the attack for reflections from a human eye, which managed 150 pt font from 3.5 metres away. Of particular interest is the bright light on the right hand side of the image on the right. This point of light was used for the PSF and subsequent deconvolution and was also taken from the eye itself, showing how other common light sources, such as desktop lamps, can aid an attacker. Furthermore, limitations of this attack were primarily derived from the equipment used, with focussing issues and the size of the image sensor in the camera the major limiters. Improvements in the equipment used may increase the power of this attack even further.

The authors also applied their advanced image deconvolution algorithms to diffuse light reflections on a wall. They found that using a privacy filter on a monitor actually improved their ability to reconstruct an image [7]. However, upper bounds calculated from their experiments suggested that the attack could not go beyond recognising rough shapes of large symbols on screen. Countermeasures suggested by the authors made mention of the possibility of using notch filters to block out specific wavelengths of emitted light from the monitor, since TFT monitors produce a very narrow spectrum, but such filters are expensive. This countermeasure may warrant some further research though, as the alternative countermeasures, such as blocking all potentially hostile line-of-sight zones to areas where sensitive information is to be displayed, even without direct line-of-sight to the screen, may be too much of logistical issue to implement effectively.

The attack makes use of powerful telescopic equipment, but, while being obtainable, such equipment may be too unwieldy to implement in certain scenarios. Other attacks have demonstrated, however, that smaller, lower quality, imaging equipment is also capable of recovering sensitive information from smaller distances and can be brought to bear to attack plaintext data input

autonomously.

3.2.4 Automated shoulder surfing

Observation of keyboard input is a potentially powerful avenue of attack. Work on automated recovery of typed text, captured by a low resolution camera with line-of-sight access to a keyboard, has shown that this is indeed feasible with low cost equipment [8]. The attack works by word prediction through keystroke recognition; analysis of a video relay can determine the area of the keyboard covered by the typist's hands and when a key is pressed and where, including the space bar, which helps differentiate when a new word is being typed. Combining this with language analysis and error correction allows the generation of multiple interpretations of the typed text, each word ranked by likelihood. While recognition rates for such an automated attack may not be good enough, around 60% of correct words made it to the top 5 most likely words, to identify exact inputs such as passwords, such information can be determined by a trained human observer, with the automated portion of this attack showing great potential for use when transcribing large portions of sensitive text, such as a document or email.

That the attack was implemented with a simple, cheap webcam shows that even consumer-grade optical equipment is capable of supporting such an attack. Higher-end cameras and long range optical devices would allow an even greater range for the same attack, perhaps also increasing accuracy of keystroke recognition through better quality imaging. Such camera equipment is easily obtainable and is capable of capturing high quality images, enough to enable reconstruction of sensitive information, or even a physical key, with line-of-sight access.

3.2.5 Long and short range photographic duplication of a physical key

Though not strictly compromising protected data, one very interesting attack, powerful through the ability to conduct it from long distances, has been performed using photography to take pictures of physical keys, and attempting to replicate them from the images [58]. This works because the notches on keys are made to manufacturer-specific depth and location codes. The authors were able to design a semi-automated method of retrieving key notch codes from popular brands of keys, and were able to do so using a camera with a telephoto lens from a distance of close to 60 metres. They also tested a mobile handset camera from a close distance to the key, and found they were able to extract key notch code information and replicate keys successfully.

This attack clearly demonstrates the continued relevance of optical surveillance tools, and their potential applications for compromising the confidentiality of protected information and the increasing security threat posed by camera phones. Whether relying on physical security, such as controlling physical and line-of-sight access, or information security, encrypting documents and emails, sensitive information must be protected from optical attacks arising from targeted surveillance.

3.3 Optical attacks and surveillance

3.3.1 Future technology for attacks

All of the above attacks have one major aspect in common, their success rate for information retrieval from a distance is limited by their equipment. If custom-made equipment for long range attacks could be developed and tested, large sensor cameras with autofocus as an example, it would be very interesting to measure the improvements over the equipment used in previous experiments.

Such expensive equipment is not needed for attacks conducted in close proximity to the target however. Mobile phone cameras are increasing in quality each year, and are not limited to taking static photos. Consumer demand is pushing mobile phone technology forward at an astonishing rate, the image quality coming from recent camera-phones is beginning to rival contemporary point-and-click cameras. On top of this, future technologies currently in development will be an even larger improvement in image quality from current camera-phones [12, 44]. If the developers of that particular technology can deliver what they claim then high quality static and video imaging surveillance capability could end up in a conveniently small, non-suspicious form, a legitimate and legal electronic device which also has the ability to transmit data, in the hands of anyone with enough money to buy one. Tools like this are sure to make close proximity shoulder surfing attacks far more powerful then they are today.

3.3.2 Potential methods of defence for today and the future

Defences against the attacks covered here present a challenge. Optical attacks against a computer are an attractive target since information flow from the monitor to the user, or from the user to the keyboard, must by necessity be in a human readable format, easily understandable to anyone viewing it.

The act of inputting something that determines the security of sensitive data, such as passwords or PINs, in plaintext form is a serious limiting factor for implementations of secure systems. Because of the ramifications of shoulder surfing attacks on these types of inputs research has gone into the development of alternative user input systems, specifically designed to be resistant to this more basic form of optical attack. These include novel methods such as selecting characters from a virtual keyboard using gaze tracking systems [56], graphical based systems with hard to detect user selection procedures [74] and PIN entry on shuffling virtual keypads [57]. However, these methods are primarily designed for password and PIN input, which is fine for an ATM transaction, but it does not protect the secrecy of data being displayed. As long as an attacker has line-of-sight access to the display they will be able to recover the information being input and output. Even without line-of-sight access, images of reflections of the display can be recovered from reflecting objects that the attacker does have line-of-sight access to.

Perhaps the most obvious defence is to avoid accessing information that the user wishes to keep secret in a place where such attacks are possible, but this presents restrictions on the user for

access to potentially vital services. Wearable "video eyewear" may be a possible solution [89], when such products get cheap enough for widespread use, but this of course also requires a mobile/wearable computing device and input peripherals. So called "privacy filters" are widely available [1, 46], which manufacturers claim to reduce the ability of others to shoulder surf by reducing the angle that a screen is readable from the side, but there is no literature available about how they fare against long range attacks, or whether screen reflections are still present and exploitable. Notch filters were mentioned by Backes et al. [6], but were deemed to be unsuitable due to cost and low availability.

Defence by blocking all potentially hostile line-of-sight to the computer and user presents a security/availability trade-off, this may indeed be fine and easy to implement when accessing sensitive information from a fixed location, such as when using a desktop PC, but when access is needed on the move, in locations where the user has little control over line-of-sight access or trust in their surroundings, optical attacks can become very difficult to defend against.

Chapter 4

Acoustic Emanations

This chapter covers attacks exploiting acoustic signals. A short introduction adds some background context, while the main content is divided into two sections; emanations from keyboards and emanations from computer processes. The entire development of acoustic keystroke recognition attacks on keyboards is covered comprehensively, and information leakage from computer processes and their potential applications is also discussed. Finally, the use of acoustic emanations for surveillance purposes is explored, with potential attack and defence methods offered.

4.1 Introduction

Virtually all mechanical processes of sufficient magnitude, and in any environment that will allow it, emit sound. Sound as a wave carries information in the form of frequency, wavelength and amplitude which can be measured by audio capturing equipment such as microphones. Traditional methods of covertly retrieving information acoustically mainly dealt with listening bugs or wire tapping to collect spoken intelligence, but it has been noted in the past that bugs could be used to exploit compromising acoustic emanations [66] and there is anecdotal evidence to suggest bugged mechanical Hagelin cipher machines allowed decryption of produced ciphertexts through knowledge of the secret key, which was gained by analysis of the acoustic emanations transmitted to analysts via the bug [93]. Acoustic intelligence is certainly of value in the modern world as well, it is grouped under the larger umbrella term of MASINT (Measurement and Signature Intelligence) [33]. For obvious reasons, next to nothing is known about the technological surveillance capabilities of the organisations that collect such intelligence. However, academic interest in the area has grown over the past decade and a number of attacks have been presented at conferences and published in journals. This work has shown that acoustic emanations from human-computer interaction and computer processes can yield compromising information if used correctly.

4.2 Emanations from human-computer interaction

The distinctive sound of someone typing is something that a lot of people immediately think of when considering how we, as humans, interact with machines. But by recovering the acoustic signals of each keystroke and applying statistical pattern and language analysis techniques it is possible to accurately reconstruct what a person is typing. The most powerful acoustic attacks presented academically have been against keyboard input, and the potential application of such an attack for capturing login details and other secret information recovery autonomously to be used in the wild is very real indeed.

4.2.1 The first acoustic attack against a keyboard

The possibility of acoustic emanations from keyboards being used to identify keystrokes was first investigated by Asonov and Agrawal in 2004 [4]. The basic premise of the paper was that computer keyboards, and telephone and ATM keypads, could have keystrokes identified based on the sound emitted when each key is struck by profiling these sounds using a neural network, effectively "training" it to recognise keystrokes. This method was chosen due to neural networks having been successfully used in audio recognition-type experiments.

To begin with, the authors had to determine how they would train the neural network to differentiate keystroke sounds [4]. This was achieved by selecting a segment, from a fast Fourier transform of the raw audio capture, of a keystroke that best represents it and labelling it with a value that represented that particular key. By visually analysing the direct frequency spectrum of a keystroke it was found that the sounds produced comprised of two distinct peaks, the first being caused by the key being pressed down and the second by its release. Since the release peak was substantially lower it was decided to use the first "push" peak. Upon further examination of the push peak, however, it was discovered that there were two further peaks, separated by a short interval. Figure 3 (left image) shows how an acoustic keystroke signal can be characterised into distinct peaks. These peaks were found to correspond to the typists finger first touching the key, a short interval, and then depressing it, causing the key to hit the keyboard supporting plate, both actions causing the plate to vibrate. It was found later that selecting one of these extra peaks caused an increase in the rate of recognition, so the authors decided to make all further representative selections from the first "touch" peak, as this was found to be better expressed in most keystrokes samples. Furthermore, each touch peak was also found to be visually distinguishable from those of other keys.

The neural network was then trained with 100 keystroke samples, touch peaks being the representative labelled feature [4]. The authors tested the network by attempting to differentiate between two keys, "k" and "l", which are adjacent to each other on a QWERTY keyboard, with audio samples recorded from a standard omni-directional microphone from a distance of less than a metre away. It was found that the neural network had a high recognition rate, with only one error every 40 keystrokes. They then decided to test the effect of distance on recognition rate. Using a parabolic microphone from a distance of 15 metres, in an environment with background noise too, they were able to achieve the same recognition rate as before. After

this, samples of multiple keys were tested from each of the following 30 keys from a QWERTY keyboard;

- From the “q” key across to the “p” key
- From the “a” key across to the “;” key
- From the “z” key across to the “/” key

These keys, chosen as representative of the text area of a QWERTY keyboard, were tested against 300 keystrokes. The results showed that the keys were definitively recognised 79% of the time, and identified 7% of the time as a second choice, and 2% as a third. In just 12% of the tests was the key not recognised.

Next, the authors decided to test the possibility of using a neural network trained on one keyboard to recognise keystrokes from two other keyboards of the same type [4]. The results of this experiment showed that the recognition rate was vastly lower than the preceding tests, with only 52% of the four guesses offered by the neural network being correct for the first new keyboard, and 50% being correct for the second new keyboard. The authors noted that these results may still be useful in password snooping, as they could significantly reduce the computational effort of a brute force attack. The last experiments with a computer keyboard by the authors involved typing with variable force and variable speed.

As the preceding experiments had been testing against one finger typing with consistent force when testing against variable force keystrokes the results of the test were predictably poor [4]. However, it was found that if the neural network was trained with variable force samples it would achieve recognition rates that were almost as good, one error every 20 keystrokes, as the consistent force rates. As for typing speed, which the authors conceded was very important in considering practical applications of this attack, an experiment was carried out with the neural network being trained on one person, and then tested against keystroke samples from three different people who were free to use any typing style on the same keyboard. The results showed that though there were more errors in key recognition, typing style only slightly affected recognition rates, and that a neural network trained on one attacker for use against other people with different typing styles could possibly be developed into a viable attack.

The authors then decided to investigate why keystrokes produced different sounds. By examining the structure of the keyboard they had been using to train the neural network, they found that removing keys from the area of tested keys, in this case “k” and “l”, did not affect recognition, and that swapping the keys around resulted in recognition of “k” as “l” and vice versa [4]. They deduced from this that sound features were not a result of the surrounding environment of the key or differences in the structure of the keys themselves. They then decided to investigate the effect of the keyboard support plate on key recognition by cutting the plate around each key so only a small section of plate remained. Subsequent tests revealed that the neural network was unable to recognise keystrokes. The authors concluded that the support plate acts as a kind of drum, resonating when a key is touched and giving out identifiable acoustic emanations depending on where on the board the key is located. As a result of this the authors suggested the following countermeasures;

- The use of silent keyboards made of rubber
- Touch-screen keyboards or virtual keyboards
- Traditional mechanical keyboards that use a support plate material that does not conduct vibrations that give out acoustic emanations
- A design that does not involve placing all the keys on one plate

Further experiments on other types of keypads yielded interesting results; notebook keyboards were found to be less vulnerable, with an error rate of two keys out of 20, and both ATM and telephone keypads were highly vulnerable with all test keystrokes being recognised [4]. It was also found that a neural network trained on one telephone keypad could be used to recognise keystrokes from other pads of the same type, but recognition rates would be lower, and could vary greatly from keypad to keypad.

4.2.2 Improvement by automation of feature recognition

Soon after the work by Asonov and Agrawal, a variation of their method was devised and implemented by Zhuang et al. [95]. The authors of this paper argued that the need to manually label training samples prior to testing was a greatly limiting factor to the seriousness of the attack, and that their method allowed, in real time, after calibration, high rates of keystroke recognition;

- 90-96% of characters
- 75-90% of words from typing in English
- 90% of random characters

These recognition rates being achieved with a statistical model trained from 10 minutes of audio taken from the target and 30 minutes of computation time on a desktop computer.

The same source of acoustic emanations chosen in the experiments of Asonov and Agrawal, that is to say the sound made when a finger touched the key, was used to determine if an improvement over the use of fast Fourier transforms for feature recognition could be found. The authors found that they could greatly improve recognition rate by using cepstrum features extracted from a larger sample, one which covers the entire push peak, produced by a finger both touching the key and the key hitting the support plate after being pushed down [95]. Figure 3 (left image) shows how much more information a signal can contain if the “hit” peak is also taken into account. This procedure could be automated by extracting the features from the first 40 milliseconds of audio capture from each signal, keystrokes automatically being detected by energy levels calculated from the sum of all fast Fourier transform coefficients. Example energy levels are shown in Figure 3 (right image), demonstrating the large spikes in acoustic energy that allow for automatic detection after the energy reaches a set threshold. Such a

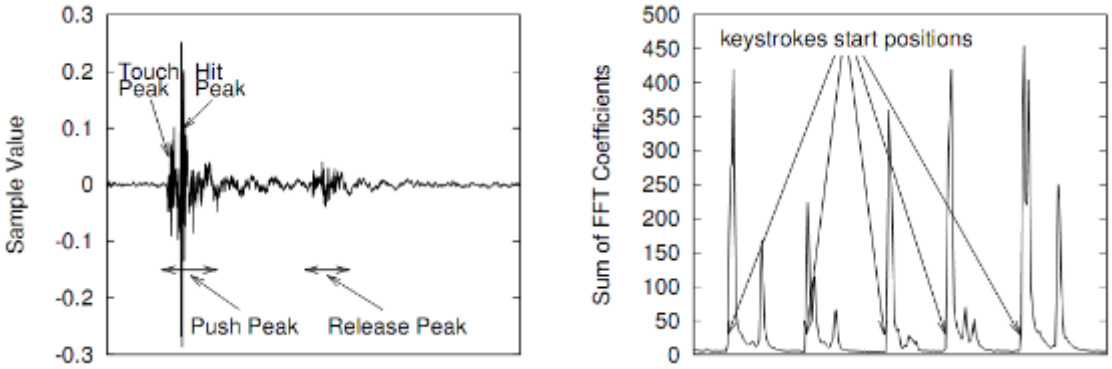

Figure 3: Representation of an acoustic signal and energy levels of keystrokes.

Figure, merged from two separate figures, otherwise unaltered, from [95]. The acoustic signal (left) is annotated to demonstrate the properties of the "push" peak, while the energy levels (right) indicate the trigger for autonomous keystroke detection.

threshold may also allow for the calibration of an attack system to optimise signal detection in noisy environments.

Next, the recognition phase; the authors clustered the recording data into 50 classes, this amount was chosen as it was greater than the 30 keys being sampled which allowed more information to be captured but also made the system more susceptible to noise [95]. To determine keystrokes by class values the authors used a hidden Markov model (HMM), with the Expectation-maximization algorithm to define class parameters, which could be improved by manually setting space bar classes, which make a distinctive sound, in the HMM output matrix. The Viterbi algorithm is then used to find the most likely sequence of keys. Manually setting space bar classes, however, is not strictly necessary as enough runs of the Expectation-maximization algorithm with random initial values can still achieve good results.

After this stage error correction is implemented. This is done by a simple spell checker and a statistical model of English grammar combined into a HMM [95]. Recognition rates can then be vastly increased using a feedback method. This is done by choosing words with less than a quarter of their characters being corrected in the previous step, since these words are more likely to be correct. They are then used as labelled samples and run through the recognition phase again, with the output going through error correction. The process can be repeated until error correction yields no significant improvement and the quality of classes can be determined by the level of error correction they require. It is important to note that when identifying random keystrokes a language error detection system is still used, but solely for the ability to improve subsequent character recognition, it is not applied when checking for passwords. The authors considered using probabilistic neural networks, linear discriminant classification and Gaussian mixtures to identify words that would be passed through the feedback stage. Their results found that using simple linear discriminant classification gave the greatest recognition rates.

The authors were then able to implement their attack model in a practical setting [95]. All experiments were conducted in lower case with all letter keys, comma, full stop, Enter and

Space being used, the authors make no mention of how far away from the keyboard their microphone was but given that Asonov and Agrawal's experiments using a simple omni-directional microphone were performed from less than a metre away, this was most likely the case in these experiments too.

The first experiment was performed to see how effective their method was against one keyboard typing different texts in different environments [95]. They recorded four sets of recordings; two in a quiet environment and two in a noisy environment, each set being typed copies of different news articles. The results of this experiment are summarised above, with all characters from each set being recognised at least 90% of the time. One thing to note is that they found that their keystroke detection algorithm sometimes failed, requiring manual correction. They did however leave one set without correction to compare with the others. Each set underwent three iterations of the feedback process. The language error detection stage was also applied and compared with recognition rates where it had not been used. This stage greatly improved recognition rates of words, on average a little less than a third better compared to no language error detection. It was also found that noisy environments negatively affected recognition rates, but increased feedback iterations reduce the recognition rate gap between environments. Interestingly, it was found that the recognition rate of the set that did not undergo manual correction was negatively affected, but still achieved a human readable level of keystroke recognition. The authors determined, by analysing differing lengths of their recorded sets, that at least 5 minutes of audio should be used to attempt keystroke recovery.

Next, the authors looked at the efficacy of their method on other models of keyboards; of differing age, brand and method of connectivity, by typing one set text, the same for each [95]. Though this experiment was affected by high levels of background noise they found that these keyboards were just as vulnerable to attack as the first. Finally, the authors tested the ability of their method to detect random passwords of varying length; 5, 8 and 10. For each of the 500 random password trial attempts lists were produced that presented all possible passwords detected by the method, the most likely being at the top of the list, allowing an attack that attempted each password in turn. On average;

- 90% of 5 character passwords were found within 20 trials.
- 77% of 8 character passwords were found within 20 trials.
- 69% of 10 character passwords were found within 20 trials.

Furthermore, 80% of 10 character passwords were found within 75 trials. This demonstrates that the attack obviously has great potential to reduce the work effort of brute forcing a password.

Suggestions by the authors for improving the attack included developing methods for recognising use of special keys, such as Shift, Backstroke and Caps Lock [95]. They also suggested that the attack could be expanded to other languages, or even to detect implementation-specific use of language such as when writing software. Further development of models for the error detection and feedback stages could also be of interest, with the authors mentioning hierarchical HMMs to consolidate the separate models. It was also suggested that advances in improving

signal-to-noise ratios in other fields of signal processing may find an application with this type of attack. As for defences against the attack, they highlighted that physical security played an important role, making sure that if one needs to use a computer for accessing or typing sensitive data, to do it in a secure environment, but that the most important lesson was that reliance on typed single factor authentication, especially when that single factor is weak such as a password, results in increased vulnerability to this type of attack.

As noted above, recognition rates were much higher when utilising language analysis and correction. Another attack exclusively uses such techniques to reconstruct likely typed words from keystrokes, achieving real time recognition from acoustic emanations.

4.2.3 Implementing a practical dictionary attack

The latest published paper to deal with acoustic keyboard emanations covers the application of statistical analysis to keystroke noise on the principle that location of a key relative to previous keystrokes leaks information via differences in sound, and the method uses this information to determine the most likely candidate for the password from a pre set dictionary of English words [14]. The authors decided to develop this method of attack after coming to the conclusion that Asonov and Agrawal's method, which needed careful labelling of audio samples, and Zhuang et al.'s method, which needed 10 minutes of audio recording and 30 minutes of computation on a PC to self-train, required too much time on the part of the attacker to be implementable in a real attack. Subsequently, they designed their method to be able to function against only a short audio recording of keystrokes and require less than 20 seconds of computation on a PC per word.

The authors used the sound of each entire keystroke, which constitutes both the push peak and the release peak, as the signal to be processed [14]. Keystroke recognition was performed much like in Zhuang et al.'s paper, with raised energy levels determined by the sum of fast Fourier transforms, except that this time two main peaks, push and release, are extracted and kept to measure similarity between keystrokes. In this attack, similarity needs to be measured in order to determine key location, and therefore identity, relative to the previous key. Similarity of acoustic features gives an idea of how close the two keys are; the higher the similarity, and the closer the keys should be. Three methods were trialled for testing similarity; cross-correlation, fast Fourier transforms and cepstrum analysis. These were rated by their "Precision" and "Recall" rates, metrics introduced by the authors to compare the classification of similarity values into predicted key location relative to the true value of the word. This is shown later by extracting "constraints", the degree to which two keys are located to each other categorised by distance, based on the similarities of their acoustic features.

In addition to these trials, the authors also experimented with how best to utilise both press and release peak similarities to make a matrix of the values, by measuring rates from five variations of similarity output; press only, release only, the minimum of the two values, the maximum, and the mean [14]. They found that the best precision and recall rates were achieved by using a cross-correlation function to produce the similarities of the press and release peaks, which were then output to a single matrix by the mean-average function. The authors then had

to find a method to infer the constraints from the similarity matrix. They apparently tried many methods, and arrived at one of their own devising, citing its ability to produce balanced precision and recall rates, called the "BestFriendsPickPolicy". The policy works by separating keystrokes into one of four classes;

- "EQ" – the same key
- "ADJ" – one of the keys adjacent
- "NEAR" – a key a maximum of two keys away
- "DIST" – any key more than two keys away

These classes were based on keystroke similarity ranked for each other keystroke in decreasing order. For example, if keystroke 1 was ranked as number one in the similarity matrix row for keystroke 2, they would be called "Friends". If keystroke 2 was ranked number one for keystroke 1 as well, they would be called "Best friends", and could be inferred to be equal to each other, EQ, since the sound would most likely be coming from the same key. For lower rankings a lower classification would be issued. The authors tested this policy across three separate keyboards and found that precision and recall rates were generally consistent.

The last stage in the attack requires the constraints to be evaluated against the attack dictionary to produce a list of possible passwords, ordered by likelihood of correctness [14]. This is done by selecting a single letter and evaluating its constraints against a table of EQ, ADJ, NEAR and DIST keys relative to the position of each key, which in this case is our chosen letter. This allows the selection of dictionary words with matching constraints as those demonstrated for the letter. To perform this operation for multiple letters a Boolean matrix function can be used to evaluate multiple constraints simultaneously.

27 words were tested, with lengths of 7 to 13 characters, each word being typed on three different keyboards and processed seven times [14]. They categorised effectiveness of each attempt by ranking the place of the true password within the list output by the attack method. Overall they found their attack placed the correct password within the top 50 potential passwords 73% of the time. The authors found two main influential factors of the success of their attack. Firstly, character repetition greatly increases ranking of the correct password, with 90% of words with two or more EQ constraints placing in the top 25. They found that this seemed to correlate with the observation that EQ constraints produced the best precision and recall rates. Secondly, that word length appeared to have a positive effect on the ranking of the password. They theorised that the greater number of constraints given by longer words could increase the chance of attack success.

The authors concluded by suggesting that their method was capable of producing an effective dictionary attack password cracker based on acoustic emanations [14]. They identified the possibility of refining the attack by improving on the statistical methods they employed, and also of incorporating inter-keystroke timing into the attack. It was also acknowledged that to make this attack stronger in the real world the ability to distinguish Shift-key keystrokes, punctuation marks and numerical digits would need to be implemented.

4.2.4 Possible future implementations and variations

Taking inspiration from the methods described above, student projects dealing with acoustic emanations of keyboards have indicated areas for possible development of further attacks.

One project [43], drawing from previous work in person-specific keystroke characteristics [13] and the acoustic keyboard work of Asonov and Agrawal, and Zhuang et al., looked into developing a small scale user authentication system based on combining inter-keystroke timing characteristics with acoustic keystroke emanations. Though not accurate enough to employ as a user authentication system, with a false accept rate of, at best, 12.9%, the ideas presented may find use in surveillance settings, in situations where it is impossible to get an image of a user or identify them by other means, and identification is required beyond the recorded use of login details.

Another project dealt with keystroke identification through triangulation via two microphones [37]. The students were able to differentiate between keystrokes some distance apart with high accuracy by clustering keystroke data by the difference in received time at each microphone. The accuracy of keystroke recognition decreased the closer the two keys were, however. This method could be viewed as a variation on the work done by Berger et al. [14], instead of identifying key location by acoustic similarities produced by key location on the plate, it is done by differences in time taken for the sound to be received by two separate microphones. This raises the possibility of a functional version of the proposed triangulation attack by using the same type of statistical language and key location techniques employed by Berger et al., applied to keystroke distance clusters instead of frequency peaks.

Attacks against keyboard input perfectly demonstrate the difficulty in protecting secret information that has to be accessed or transcribed using plaintext data input. This factor readily lends itself to surveillance techniques, as will be discussed later.

However, keyboards are not the sole focus of acoustic attacks, with some work also being done on emanations from PCs.

4.3 Emanations from computer processes

To a lesser extent, acoustic emanations from computational processes can also reveal information to an attacker. Though it is generally limited to identifying computationally intensive activity and processor instructions, it is possible to use known acoustic signatures as a covert channel for the relay of small amounts of sensitive data.

4.3.1 Acoustic signatures from RSA signatures and processor instructions

Firstly, Shamir and Tromer have shown how, as a proof-of-concept presentation, computer processing components can leak basic information through acoustic emanations [76]. By measuring acoustic signals from an open PC case with the microphone 20 cm away, they were able

to identify x86 "HLT" instructions through analysis of a spectrogram of the captured audio. By solely running GnuPG, specifically an RSA signing operation with a large key, they were able to identify when signing started and stopped, and even identify the two exponentiation stages of RSA using CRT. They also found that different RSA keys produced different audio signals, with distinctive emanations for each modular exponentiation, when signing the same message. When investigating the source of these emanations the authors found that freezing the capacitors on the PCs motherboard during another type of CPU instruction produced a great deal of noise in the signal for the rest of that instruction, indicating that these capacitors were responsible for the distinctive acoustic emanations.

4.3.2 The potential for a covert channel

Continuing on from the work by Shamir and Tromer, some related work has been published replicating and extending the attack, investigating the feasibility of using acoustic emanations from a practical attack perspective [59]. The authors were first able to replicate Shamir and Tromer's findings with respect to differences in RSA audio signals from capacitors on the motherboard of the PC being tested. By using a more powerful modern processor they were able to select the clock speed for the operation of test CPU instructions. They found that instructions at higher CPU speeds produced acoustic emanations of a lower amplitude, up to 2.4 GHz, where no signal was detected, which the authors conceded may be a result of the shortcomings of the equipment used. The authors reasoned that this confirmed the observation that the source of the emanations were indeed the capacitors, and that higher clock speeds caused increased voltage across them which lead to physical movement between the capacitor plates, and thus an acoustic signal. The authors also investigated the possibility that acoustic signals could be used as a covert channel for data transmission in the form of distinctive emanations from four processor instructions. These four signals could be interpreted as data in quaternary form, an extension of the possibility of using two signals to transmit binary data. In an experiment, the authors used a keylogging program to transmit quaternary data at 2 seconds each keystroke, using a wireless camera/microphone fitted into the PC case. They noted that the low data rate and conspicuous noise of this method mean that it would not be suited to real time eavesdropping, but should be useful for relaying specific captured keystrokes, such as login details, to be captured remotely.

From the published work available it seems that acoustic emanations from processors have yet to constitute a viable threat against cryptographic implementations. The most practical attacks from a surveillance point of view would seem to be those demonstrated against keyboards, and there are interesting ways they can potentially be carried out.

4.4 Acoustic emanations and surveillance

4.4.1 Potential methods of attack

The attacks discussed above are currently the most developed of their kind that have been published openly. Concrete evidence of modern attacks on security systems through the use of acoustic emanations is few and far between. There have been a couple of cases reported in the media though, news of UN officials having conversations being bugged through mobile phones [10] and of the same technology being used against members of organised crime gangs [21]. This type of technology is also available commercially [39]. However, this is merely surveillance of conversations. One interesting avenue of attack could be to combine the surveillance capabilities of a mobile handset with the acoustic attack methodologies outlined above, using the handset's microphone to relay acoustic signals back to a receiver. The signals can then be taken and stored for later analysis, or passed on to a PC or laptop and processed in real-time, using improved implementations of the signal processing and analytical methods developed above, making for a potent remote attack. The ubiquity of mobile handsets, the lack of suspicion they arouse and their increasing build quality and data transfer capabilities make them a potentially excellent method of collecting acoustic emanations covertly, either by maliciously infecting a target's own handset, or placing one near the target's computer. Previous work on the audio capture capabilities of handsets combined with methods for environment classification [62] could be developed into advanced automation of signal filtering for improved clarity of acoustic emanations captured by the types of microphones used in handsets. Laptop microphones are also a promising target, as has been demonstrated [36], but they may be redundant for direct use against a target, since the laptop would need to already be compromised with malicious software. The wireless microphone used by Lemay and Tan [59], or the parabolic microphone used by Asonov and Agrawal [4] are other alternatives to these types of technologies, sometimes sold as hobby "spy kits", they are well within the ability of any determined attacker to implement.

The limitations of acoustic attacks must be underlined however, the most effective attack investigated utilises statistical properties of the English language to generate possible lists of words typed. This is fine for predicting dictionary passwords, but these are weak forms of authentication anyway, and over the past years great efforts have been made by security experts to convince users not to use bad passwords. Any security savvy target worth his or her salt will most likely already use strong alphanumeric and symbol passwords, for which the properties of the English language cannot be used to predict the content of. If access to a secured system is not desired though, the attack would be an excellent way of determining the content of outgoing encrypted emails and typed documents. The high random character recognition rates achieved by Zhuang et al. required previous audio capture and computation time of data from the target [95], but if this can be achieved before input of the desired password even a "strong" alphanumeric password can be determined with high accuracy.

4.4.2 Potential methods of defence

Coming up with adequate countermeasures against these types of attack may not be so simple. Changes to keyboard design would only be effective for an individual or an organisation with the funds to implement custom made equipment. This countermeasure is unlikely to find its way to the general public unless keyboard manufacturers are mandated by law or security standards, or find some economic reason to improve the design of their keyboards. Accessing sensitive systems and typing sensitive documents within a special room designed to eliminate acoustic attack capabilities could be effective, but would most likely turn into a logistical nightmare. Doing the same in an environment surrounded by random noise of a higher amplitude then the emanations produced may introduce noise to make the acoustic emanations unrecoverable, but would be distracting to the user, and if that environment were a public place it would likely open up an avenue for optical attacks.

These attacks show great potential to become more powerful in the years to come, especially when taking into consideration the growing capabilities of mobile computing platforms such as smartphones, and their development may yield some exciting developments in the security of human-computer interaction, as traditional methods, such as dedicated audio bugs, become less suited to emerging threats.

Chapter 5

Discussion

Having reviewed the development of published attacks utilising exploiting compromising emanations, it is important to put this knowledge into context with their potential application for surveillance in the real world. This chapter intends to do just that, and addresses the aims set out in the introduction of this project. Firstly, the applicability and limitations of these attacks are considered, and their practicality assessed. They are related to current security trends and possible implementations of emanation surveillance, using various sensor technologies, are discussed. Finally, new directions in the research and development of signal processing and human-computer interaction methods, which may in future be relevant to these attacks, are examined, and possible security implications of surveillance technology deployment are explored.

5.1 Current considerations

5.1.1 The current status of these attacks and their technical limitations

The compromising emanations reviewed here have obvious applications for immediate use in conjunction with surveillance tools. Their power stems from the exploitation of ubiquitous methods of data input and output, namely keyboard and information display use. The sources of these emanations are common across such devices; LCD monitor video cables and keyboard design are generally the same across makes and models, which increases the applicability of these attacks. Furthermore, the requirement of input and output data, using keyboards and displays, to be human readable, and therefore also understandable to an eavesdropper, bypasses the security offered by implementing cryptographic data protection measures.

To a certain extent the attacks are limited by the quality of the equipment employed. For instance, an attacker will have to spend a lot of time optimising custom assembled radio receiving equipment for EM eavesdropping, as opposed to using a receiver specifically designed to test

TEMPEST standards. Optical attacks require equipment of reasonable quality to generate readable images of screens, or need to be small and reliable enough to relay video feeds of keyboard input from hidden locations. With regards to acoustic emanations, while none of the microphones used in published attacks were especially expensive or difficult to obtain, the data processing requirements of accurate keystroke recognition mean that the attack is far from portable, requiring either a method of acoustic data relay with a high enough bit rate, or the use of parabolic microphones, to fully realise a remote attack.

Additionally, all the attacks presented in this project require captured raw data to have a good signal-to-noise ratio. Noisy measurements are capable of drastically reducing the capabilities of both autonomous systems and human operators to correctly recover any desired information from the compromising emanation.

5.1.2 Are these attacks practical, or just theoretical?

Considering the issue of practicality; for any determined attacker, the tools needed for the successful carrying out of these attacks, at their most basic level, are readily available and simple to implement. Telephoto and zoom lenses and good quality camera equipment, radio frequency receivers and open source software radio tools, and embedded microphones in conjunction with signal processing and statistical analysis tools all allow the implementation of EM, optical and acoustic attacks.

Further increasing the threat that compromising emanations pose is the fact that adequate defences are hard to implement in a convenient manner. For a device being used in a static location within an organisation, for instance, securing line-of-sight access, EM shielding of components and perhaps the room, and sweeping for bugs and prevention of employees using personal electronic devices that could be used as such, would be fairly straightforward to implement, but time consuming and onerous for all involved, physically restricting data access that may previously have been allowed over a local network.

Any individual could also attempt similar measures, solely accessing or creating sensitive data in private, away from uncovered windows, using a laptop on battery power to avoid power line emanations, using a keyboard that limits EM and acoustic emanations, or if suspecting compromisation of their own property, operating their computers in the presence of high amounts of background noise. However, the level of security provided by technical and monetary backing and physical security measures, as is found within an organisation, will still elude the individual attempting to secure themselves against these types of threats.

Furthermore, when access to sensitive data is required on the move, effective defences become more difficult to deploy. In public places the target has little to zero control over line-of-sight access to a display, if a more private location is sought, they still cannot trust that this location is secure from other forms of attack such as audio or video bugs and EM monitoring equipment.

These factors together signify that attacks exploiting compromising emanations, using published methods, are now, even at an early stage of development, practically implementable in certain

scenarios. If real world attacks using more basic forms of surveillance, such as covert cameras in ATMs [9], or the bugging of conversations for intelligence gathering [10, 21], are anything to go by, attacks based on the compromising emanations discussed here may be implemented in the future, when simple methods of bypassing confidentiality offered by common cryptographic implementations, such as the encryption of emails and the input of passwords as authentication for whole disc encryption, may be desired.

5.1.3 How does this relate to the wider security landscape?

The proliferation of e-commerce and electronic data storage and access services have resulted in an increased need for information security, and the security services provided by cryptographic primitives such as symmetric and asymmetric ciphers. Since most openly published and standardised ciphers tend to be well designed, and able to resist logical cryptanalysis, the majority of viable real world attacks against these ciphers will continue to be those that exploit side channels. In turn, this results in the improvement of existing technology and the development of side channel resistant cryptographic implementations.

Indeed, a lot of the research already undertaken within the field of side channel analysis has been in the recently expanding area of smart cards. The push to roll out such technologies, and other embedded systems, has resulted in the further implementation of cryptography within consumer grade electronics. While cutting edge cryptographic methods are applied to these new technologies, confidentiality can still be broken on the majority of current computerised systems because of the necessity of plaintext human-computer interaction, and the information leakage that can be exploited through targeted surveillance of such interaction. Breaking the confidentiality of the sensitive data required for authentication and secure data access, such as typed passwords, can therefore become a realistic aim of any determined attacker.

Traditionally, surveillance has mainly involved the monitoring of people and their actions, gathering intelligence through physically observing comings and goings, and at which times and locations. This type of surveillance, such as observation with basic optical equipment, can be used preliminarily to identify any vulnerabilities the target may possess, or investigate potential avenues for other types of attack. In this sense, it is possible to regard any information about the target recovered covertly as useful. Indeed, by building up as detailed as possible an assessment of the target, an attacker can maximise the effectiveness of any further attacks they might employ. In fact, methods of doing so are applicable for the preparation of a variety of established attack methods, not just those of a passive nature.

For instance, traffic analysis of telecommunications data, provisioned for law enforcement and intelligence agencies under RIPA [67], and also implementable by attackers using software and hardware network sniffers, can be used to infer information about the target and their capabilities and occurrences of the transmission of encrypted communications can indicate typical times and methods of communication that are likely to be of interest. Social engineering, where an attacker can persuade, coerce or trick a target into revealing relevant information about themselves [34] is another viable method of sizing up a target's vulnerabilities.

The implementation of attacks based on compromising emanations is likely to simply be an extension of the types of attacks we already see occurring in the wild; it has already been demonstrated how they can drastically reduce the work effort of a brute force attack on even a "strong" password. Any increase in pace of the development of technologies and methods to exploit and defend against compromising emanations, both in research and in the wider world, will be tied closely to the success of cryptographic implementations and their resistance to side channel cryptanalysis. At a time when computer crime is on the increase, cryptographic implementations are beginning to become more visibly implemented, and with greater emphasis being placed on security in software design, attacks following the path of least resistance, plaintext human-computer interaction, will only become more common.

5.1.4 How can these attacks be implemented?

It has been demonstrated that the exploitation of compromising emanations can be used to bypass the confidentiality offered by cryptographic primitives by targeting ubiquitous methods of plaintext data input and output, with the most vulnerable security measures being passwords. Furthermore, proven methods using easily available technology are completely passive and can be conducted remotely, highlighting the applicability of such attacks in a surveillance context, and their implementability by any determined attacker.

Of course, it is impossible to predict exactly how real world attacks utilising compromising emanations will be implemented, if any recorded attacks take place at all. However, it is possible to suggest new attack techniques by combining proven methods and technologies with newer developments in the fields of sensor technology and signal processing.

The current capabilities of cutting edge technologies allow ELINT, SIGINT and MASINT sensors to pick up communications signals and electronic, acoustic and other types of signals and relay them on for further analysis. These sensors are employed not just in static locations but on mobile surveillance aircraft to maintain a steady stream of intelligence used for assessing potential targets [92]. The application of such sensors for intelligence gathering is an area of research receiving increasing amounts of government funding, mainly for military applications, from calls for more sensitive EM sensors [23, 29, 30], to autonomous target recognition from optical sensors and camera networks [25-27]. Since the details of these projects and their exact applications are restricted, the small amount of information on them can only be used to gain a rough idea about the direction research and development of sensors for military intelligence gathering is taking.

However, it is still possible to discuss possible methods of attacks from looking at openly published research on sensor technology and the capabilities of consumer electronics. For instance, the increase in malicious software surveillance tools can open up potential avenues for attack. Mobile phones and laptops commonly have integrated cameras and microphones that could be activated for surveillance purposes [36, 91], perhaps making use of compromising emanations such as optical reflections, direct optical recognition of keystrokes or acoustic keystroke analysis. They also have wireless data transfer capabilities, allowing the remote execution of such attacks by relaying raw signal captures back to an attacker's own device for analysis [20].

Smart phones present an interesting dilemma; their increasing power, and the move towards common operating systems for use with them, is cementing their popularity among consumers, with access to emails and documents a major selling point. However, it also opens up new attack vectors, firstly with the phones themselves, which can be infected with malware and used as malicious surveillance tools by an attacker [17, 39, 85], and secondly by allowing on the move access to potentially sensitive data, increasing opportunities for surveillance and other attacks. The attraction that smart phones possess as a target for attack has already drawn attention from not just criminals, but government organisations, ranging from clandestine attempts to mass infect smart phones at the distributor level [83], to official government mandates cutting off smart phone services, as a ploy for coercing phone manufacturers to gain access to encrypted networks owned by them [11, 84]. As yet, there is no published work dealing with compromising emanations from small devices such as smart phones, but there may be in the near future, since these types of attacks present an attractive method for defeating any implemented data protection measures passively.

5.2 Considerations for the future

5.2.1 Future technology and future research

With a great deal of research being conducted on improving sensor technology, EM, optical and acoustic attacks are likely to be the first types of attacks seen in the wild. Both government, military and consumer electronics research and development are likely to push forward the further miniaturisation and data capture capabilities of newer, better sensors.

Research into the consolidation of data received from multiple sensors of different signal sources for improved data acquisition will no doubt find an application in new surveillance technologies. Systems already implementing such techniques include a maritime security surveillance network [31], a surveillance robot designed to assist the housebound in case of emergency [94], a car park surveillance system using both audio and video feeds to identify possible security incidents [64], and the use of multiple video feeds to automatically extract biometric human gait characteristics [42], and techniques are being researched that allow the use of multiple cameras to produce a static image with a vastly increased resolution compared to the resolution of each individual camera [28]. The issue of multiple available sensors is also highly relevant to the potential for smart phones to function as surveillance tools, research has already been published about the potential commercial aspects of a smart phone with the capability to accurately assess its locality through information from its built in sensors [5].

This research may be used as the basis for developing an attack using all possible sources of compromising emanations to accurately reconstruct data. By using multiple sources, a higher signal-to-noise ratio could be achieved, which could allow an attack to be conducted at a greater range. It would also allow successful information retrieval from a system only partially defended from emitting compromising emanations, as long as one viable source of emanations remained, an attack would still be successful.

One other factor to consider is that both traditional surveillance and attacks utilising compromising emanations require a human operator to be present, and to guarantee the recovery of useful information an attacker would have to be able to monitor the target constantly. While a certain degree of autonomy has been demonstrated in certain attacks, such as the automatic recognition of keystrokes, true autonomy in the future would allow the deployment of drones or networked sensor stations working and relaying data without the need for human interaction, save a general system for command and control.

Obviously, such technology is beyond the resources of an everyday attacker. However, as demonstrated in published attacks, techniques for manual control of data recovery, using cheap, readily available equipment, are feasible for any determined attacker to implement. The additional exploitation of common, consumer electronics, such as smartphones and laptops and their computational potential, can increase the power of surveillance attacks, especially if the statistical analysis techniques required for attack autonomy were made widely available as software code, allowing anyone with the appropriate equipment to implement an attack easily. For example, research into the development of automatic speech recognition and translation systems, using statistical models similar to those implemented for autonomous keystrokes detection i.e. hidden Markov models and Gaussian processes, and improved signal processing from multimedia sensors, for use on platforms with lower computational power such as mobile devices [2, 80, 96] could be applied to surveillance technology to further improve attack capabilities.

The development of new methods of human-computer interaction will also be an important factor in determining the relevancy of the attacks covered here. Since the power of these attacks relies on plaintext leakage from input and output devices, research on new methods such as gaze based PIN entry [56], or even using a human brain-computer interface to authenticate to system using thoughts [87], could not just revolutionise how users interact with computing systems, but could preclude some existing methods of exploiting compromising emanations and open up avenues for new ones.

As noted earlier, the development and implementation of countermeasures built into existing, vulnerable, devices will depend on how widespread and high profile attacks using compromising emanations become. Research into low-cost and easily implemented defensive measures is likely to be the most relevant to current attack capabilities at the moment; since current defence strategies within private organisations rely on projected risk management, their main focus will be on defending against the much more numerous network based attacks and exploits taking advantage of software vulnerabilities. Government and military organisations already have their own, confidential, standards dealing with the mitigation of compromising emanations. Individuals hoping to defend themselves can only attempt to implement suggested countermeasures suggested by the authors of each published attack, since, at this point in time, these are the only suggested defence measures available.

5.2.2 The security of implementations

One other aspect of interest is that of the security of sensor networks themselves. Some of the published attacks covered here have demonstrated that covert channels can be used to

relay recovered data covertly and, potentially, in an encrypted format [55, 59, 61], to prevent discovery by the target. This idea can be extended to sensor networks themselves, as their design does not usually incorporate protection against traditional side channel attacks [70], in fact, remote, comprehensive diagnostic analysis by power consumption has been shown to be possible on sensors, underlining just how much information can leak from such devices [47]. Research into the implementation of cryptographic primitives across sensor nodes to protect the confidentiality of recovered data may be applicable to protect relayed data [81].

In fact, confidentiality may not be the only security service that might need to be employed in a sensor network. Integrity of the data would be of the utmost importance, since corrupted data would prevent the successful retrieval of the desired data. Of even more interest would be that of maintaining availability. The old idiom; “The best defence is a good offence”, may find an application in protecting against the eavesdropping of compromising emanations, a target concerned about the possibility of such attacks could purposefully flood the area around his own devices, or even target a suspected sensor, with noise; loud, random audio signals, bright flashing lights and random RF interference. If all properly calibrated, this would deny an attacker both integrity, and perhaps even completely blocking availability, of acquired data from their sensors, indeed, highly sensitive sensors may be even be damaged by exposure to high levels of noisy data.

Chapter 6

Conclusion

The monitoring of side channel leakage from human-computer interaction has been demonstrated to allow the systematic, targeted, surveillance of compromising emanations, which offers a remote and passive method of breaking the confidentiality provided by cryptography. EM, optical and acoustic emanations have all been shown to leak enough recoverable information to allow this. By recovering leaked plaintext data from a target, an attacker can bypass data confidentiality mechanisms. Plaintext leakage from human-computer interaction is so revealing that even if the attacker did not immediately understand the data recovered, due to it perhaps being in another language, or some code or format suitable for the human brain to process without computational aid, it would be trivial to recover the desired information with sufficient resources.

It has also been discussed here how cutting edge sensor technology and analytical techniques will continue to increase the power of attacks utilising compromising emanations and their applicability for surveillance operations. While the research and development of technology and methods of implementation for surveillance by military, intelligence and law enforcement organisations is generally classified, and beyond the scrutiny of an open peer review system, published advances in commercial sensor technology and developments in data input and output technologies will continue to contribute towards the development of attacks and defences based on the exploitation of compromising emanations.

Possible future attacks utilising such emanations in the wild may force greater attention to be paid towards the greater implementation of technical countermeasures and risk mitigation strategies. Although these types of attacks seem to have been known in military and intelligence circles for some time and the cutting edge technology will most likely remain under wraps, the open publication of methods for exploiting emanations will continue to push forward the development of new technologies and approaches for attack and defence, benefiting not just those working within security research and industry, but everyday consumers and users too.

Bibliography

[1] 3M, Privacy Filters: 3M UK & Ireland. Available at: http://solutions.3m.co.uk/wps/portal/3M/en_GB/vikuiti-uk/home/products-solutions/privacy-filters/ [Accessed July 4, 2010].

[2] Aminzadeh, A.R. & Shen, W., 2008. Low-resource speech translation of Urdu to English using semi-supervised part-of-speech tagging and transliteration. In *Proceedings of the Spoken Language Technology Workshop, 2008.* pp. 265 – 268.

[3] Anderson, R., 1993. Why cryptosystems fail. In *Proceedings of the 1st ACM conference on Computer and communications security.* pp. 215-227.

[4] Asonov, D. & Agrawal, R., 2004. Keyboard acoustic emanations. In *Proceedings of the IEEE Symposium on Security and Privacy, 2004.* pp. 3-11.

[5] Azizyan, M., Constandache, I. & Choudhury, R., 2009. SurroundSense. In *Proceedings of the 15th annual international conference on Mobile computing and networking - MobiCom '09.* pp. 261-272.

[6] Backes, M., Dürmuth, M. & Unruh, D., 2008. Compromising Reflections-or-How to Read LCD Monitors around the Corner. In *Proceedings of the 2008 IEEE Symposium on Security and Privacy.* pp. 158-169.

[7] Backes, M. et al., 2009. Tempest in a Teapot: Compromising Reflections Revisited. In *Proceedings of the 2009 IEEE Symposium on Security and Privacy.* pp. 315-327.

[8] Balzarotti, D., Cova, M. & Vigna, G., 2008. ClearShot: Eavesdropping on Keyboard Input from Video. In *Proceedings of the 2008 IEEE Symposium on Security and Privacy.* pp. 170-183.

[9] BBC, 2003. Gangs preying on cash machines. Available at: http://news.bbc.co.uk/1/hi/uk/3157214.stm [Accessed July 4, 2010].

[10] BBC, 2004. 'This goes no further...'. Available at: http://news.bbc.co.uk/1/hi/uke/3522137.stm [Accessed July 3, 2010].

[11] BBC, 2010. Gulf states unveil Blackberry ban. Available at: http://www.bbc.co.uk/news/world-middle-east-10830485 [Accessed August 3, 2010].

[12] BBC, 2010. Quantum tech boosts phone cameras. Available at: http://news.bbc.co.uk/1/hi/technology/8580372.stm [Accessed June 3, 2010].

[13] Bergadano, F., Gunetti, D. & Picardi, C., 2002. User authentication through keystroke dynamics. *ACM Transactions on Information and System Security*, 5(4), 367-397.

[14] Berger, Y., Wool, A. & Yeredor, A., 2006. Dictionary attacks using keyboard acoustic emanations. In *Proceedings of the 13th ACM conference on Computer and communications security*. pp. 245-254.

[15] Bernstein, D., 2005. Cache-timing attacks on AES. Available at: http://cr.yp.to/antiforgery/cachetiming-20050414.pdf [Accessed June 22, 2010].

[16] Betts, D.E., 1970. *Working Against the Tide (COMSEC Monitoring and Analysis), Parts 1 and 2*, National Security Agency.

[17] Bickford, J. et al., 2010. Rootkits on smart phones: attacks, implications and opportunities. In *Proceedings of the Eleventh Workshop on Mobile Computing Systems & Applications*. pp. 49-54.

[18] Bogdanov, A., 2008. Multiple-Differential Side-Channel Collision Attacks on AES. In *Cryptographic Hardware and Embedded Systems – CHES 2008*. pp. 30-44.

[19] Brumley, D. & Boneh, D., 2005. Remote timing attacks are practical. *Computer Networks*, 48(5), 701-716.

[20] Cai, L., Machiraju, S. & Chen, H., 2009. Defending against sensor-sniffing attacks on mobile phones. In *Proceedings of the 1st ACM workshop on Networking, systems, and applications for mobile handhelds - MobiHeld '09*. pp. 31-36.

[21] CNET News, 2006. FBI taps cell phone mic as eavesdropping tool. Available at: http://news.cnet.com/FBI-taps-cell-phone-mic-as-eavesdropping-tool/2100-1029_3-6140191.html [Accessed July 3, 2010].

[22] Cryptek, Emanation Security & TEMPEST Products. Available at: http://www.cryptek.com/products/emanation_security_products [Accessed July 20, 2010].

[23] DARPA, ADR - Programs - Microsystems Technology Office. Available at: http://www.darpa.mil/MTO/Programs/adrt/index.html [Accessed July 20, 2010].

[24] DARPA, Channelized SIGINT/ELINT Receiver for UAV Apps - Microsystems Technology Office. Available at: http://www.darpa.mil/MTO/Programs/chaser/index.html [Accessed July 20, 2010].

[25] DARPA, IPTO :: Programs :: ARGUS-IS. Available at: http://www.darpa.mil/ipto/programs/argus/argus.asp [Accessed July 20, 2010].

[26] DARPA, IPTO :: Programs :: VIRAT. Available at: http://www.darpa.mil/ipto/programs/virat/virat.asp [Accessed July 20, 2010].

[27] DARPA, IPTO :: Programs :: VIVID. Available at: http://www.darpa.mil/ipto/programs/vivid/vivid.asp [Accessed July 20, 2010].

[28] DARPA, MONTAGE - Programs - Microsystems Technology Office. Available at: http://www.darpa.mil/MTO/Programs/montage/index.html [Accessed July 20, 2010].

[29] DARPA, RADER - Programs - Microsystems Technology Office. Available at: http://www.darpa.mil/MTO/Programs/rader/index.html [Accessed July 20, 2010].

[30] DARPA, STO: Semiconductor-Tuned High-Temperature Superconducting Filters for Ultra-Sensitive RF Receivers. Available at: http://www.darpa.mil/sto/programs/surf/index.html [Accessed July 20, 2010].

[31] DeWeert, M.J., 2009. Fusion of Information from Disparate Electro-Optical Imagers for Maritime Surveillance. In *Harbour Protection Through Data Fusion Technologies*. pp. 209-221.

[32] Dhem, J. et al., 1998. A Practical Implementation of the Timing Attack. In *Proceedings of the The International Conference on Smart Card Research and Applications*. pp. 167-182.

[33] DIANE Publishing Company, 1996. *Intelligence Threat Handbook*, DIANE Publishing.

[34] Dittrich, D. & Himma, K.E., 2006. Hackers, Crackers, and Computer Criminals. In *Handbook of Information Security*. Wiley, Vol. 2, pp. 154-171.

[35] van Eck, W., 1985. Electromagnetic radiation from video display units: an eavesdropping risk? *Computers and Security*, 4(4), 269-286.

[36] Farley, R. & Wang, X., 2009. Roving Bugnet: Distributed Surveillance Threat and Mitigation. In *Emerging Challenges for Security, Privacy and Trust*. pp. 39-50.

[37] Fiona, A.H.Y. & Pak, H.S., 2006. *ERG4920CM Thesis II - Keyboard Acoustic Triangulation Attack*. The Chinese University of Hong Kong.

[38] Fischer, W. et al., 2006. Differential Power Analysis of Stream Ciphers. In *Topics in Cryptology – CT-RSA 2007*. pp. 257-270.

[39] Flexispy, Catch Cheating Spouses with FlexiSPY - Spy Phone, GPS Tracker, Location Tracking, Remote Listening for Mobile / Cell Phones. Available at: http://flexispy.com/ [Accessed July 3, 2010].

[40] Gandolfi, K., Mourtel, C. & Olivier, F., 2001. Electromagnetic Analysis: Concrete Results. In *Proceedings of the Third International Workshop on Cryptographic Hardware and Embedded Systems*. pp. 251-261.

[41] Gebotys, C.H., Ho, S. & Tiu, C., 2005. EM Analysis of Rijndael and ECC on a Wireless Java-Based PDA. In *Cryptographic Hardware and Embedded Systems – CHES 2005*. pp. 250-264.

[42] Goffredo, M. et al., 2009. Performance analysis for automated gait extraction and recognition in multi-camera surveillance. *Multimedia Tools and Applications*, 50(1), 75-94.

[43] Gould, S., 2005. *CS229 Final Project - A Novel Approach to User Authentication Through Machine Learning of Keyboard Acoustic Emanations*. Stanford University.

[44] InVisage Inc., InVisage. Available at: http://www.invisageinc.com/ [Accessed June 3, 2010].

[45] Kasper, T., Oswald, D. & Paar, C., 2009. EM Side-Channel Attacks on Commercial Contactless Smartcards Using Low-Cost Equipment. In *Information Security Applications*. pp. 79-93.

[46] Kensington, Privacy Filters. Available at: http://us.kensington.com/html/1430.html [Accessed July 4, 2010].

[47] Khan, M.M.H. et al., 2010. Diagnostic powertracing for sensor node failure analysis. In *Proceedings of the 9th ACM/IEEE International Conference on Information Processing in Sensor Networks*. pp. 117-128.

[48] Kocher, P.C., 1996. Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. In *Proceedings of the 16th Annual International Cryptology Conference on Advances in Cryptology*. pp. 104-113.

[49] Kocher, P.C., Jaffe, J. & Jun, B., 1999. Differential Power Analysis. In *Proceedings of the 19th Annual International Cryptology Conference on Advances in Cryptology*. pp. 388-397.

[50] Koeune, F. & Quisquater, J., 1999. *A timing attack against Rijndael*, Université catholique de Louvain.

[51] Kuhn, M.G., 2002. Optical Time-Domain Eavesdropping Risks of CRT Displays. In *Proceedings of the 2002 IEEE Symposium on Security and Privacy*. pp. 3-18.

[52] Kuhn, M.G., 2003. *Compromising emanations: eavesdropping risks of computer displays - Technical Report UCAM-CL-TR-577*. University of Cambridge.

[53] Kuhn, M.G., 2004. Electromagnetic Eavesdropping Risks of Flat-Panel Displays. In *Privacy Enhancing Technologies*. pp. 88-107.

[54] Kuhn, M.G., 2005. Security Limits for Compromising Emanations. In *Cryptographic Hardware and Embedded Systems – CHES 2005*. pp. 265-279.

[55] Kuhn, M.G. & Anderson, R.J., 1998. Soft Tempest: Hidden Data Transmission Using Electromagnetic Emanations. In *Proceedings of the Second International Workshop on Information Hiding*. pp. 124-142.

[56] Kumar, M. et al., 2007. Reducing shoulder-surfing by using gaze-based password entry. In *Proceedings of the 3rd symposium on Usable privacy and security*. pp. 13-19.

[57] Kwon, T., Lee, J. & Song, J., 2009. On the Privacy-Preserving HCI Issues. In *Universal Access in Human-Computer Interaction. Addressing Diversity*. pp. 544-549.

[58] Laxton, B., Wang, K. & Savage, S., 2008. Reconsidering physical key secrecy: Teleduplication via optical decoding. In *Proceedings of the 15th ACM conference on Computer and communications security*. pp. 469-478.

[59] Lemay, M. & Tan, J., 2006. Acoustic Surveillance of Physically Unmodified PCs. In *Proceedings of the 2006 International Conference on Security & Management*. pp. 328-334.

[60] Lockton, V. & Rosenberg, R.S., 2005. RFID: The Next Serious Threat to Privacy. *Ethics and Information Technology*, 7(4), 221-231.

[61] Loughry, J. & Umphress, D.A., 2002. Information leakage from optical emanations. *ACM Transactions on Information and System Security*, 5(3), 262-289.

[62] Ma, L., Milner, B. & Smith, D., 2006. Acoustic environment classification. *ACM Transactions on Speech and Language Processing*, 3(2), 1-22.

[63] McEvoy, R. et al., 2007. Differential Power Analysis of HMAC Based on SHA-2, and Countermeasures. In *Information Security Applications*. pp. 317-332.

[64] Na, K., Kim, Y. & Cha, H., 2009. Acoustic Sensor Network-Based Parking Lot Surveillance System. In *Wireless Sensor Networks*. pp. 247-262.

[65] NIST, Cryptgraphic Modules Validation Lists. Available at: http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/1401valold.htm [Accessed July 4, 2010].

[66] National Security Agency, 1972. TEMPEST: A Signal Problem. *Cryptologic Spectrum*.

[67] Office of Public Sector Information & United Kingdom Parliament, Regulation of Investigatory Powers Act 2000. Available at: http://www.opsi.gov.uk/acts/acts2000/ukpga_20000023_en_1 [Accessed August 4, 2010].

[68] Paavilainen, R., 2002. Method and device for signal protection. Available at: http://www.freepatentsonline.com/7356626.html. [Accessed July 16, 2010].

[69] PGP, PGP 6.0.2 User's Guide in English. Available at: http://www.pgpi.org/doc/guide/6.0/en/ [Accessed July 21, 2010].

[70] Pongaliur, K. et al., 2008. Securing Sensor Nodes Against Side Channel Attacks. In *Proceedings of the 2008 11th IEEE High Assurance Systems Engineering Symposium*. pp. 353-361.

[71] Quisquater, J. & Samyde, D., 2001. ElectroMagnetic Analysis (EMA): Measures and Counter-Measures for Smart Cards. In *Proceedings of the International Conference on Research in Smart Cards: Smart Card Programming and Security*. pp. 200-210.

[72] Remote-Exploit.org, 2010. Keykeriki v2 – CanSecWest 2010 Release. Available at: http://www.remote-exploit.org/?p=483 [Accessed July 12, 2010].

[73] Rohatgi, P., 2006. Side-Channel Attacks. In *Handbook of Information Security*. Wiley, Vol. 3, pp. 241-259.

[74] Sasamoto, H., Christin, N. & Hayashi, E., 2008. Undercover: authentication usable in front of prying eyes. In *Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*. pp. 183-192.

[75] Schindler, W., 2000. A Timing Attack against RSA with the Chinese Remainder Theorem. In *Proceedings of the Second International Workshop on Cryptographic Hardware and Embedded Systems*. pp. 109-124.

[76] Shamir, A. & Tromer, E., 2004. Acoustic cryptanalysis: On nosy people and noisy machines. Available at: http://people.csail.mit.edu/tromer/acoustic/. [Accessed June 22, 2010].

[77] Siemens, 2010. Tempest Products SITEMP. Available at: http://www.automation.siemens.com/mcms/topics/en/tempest-products/Pages/home.aspx [Accessed July 20, 2010].

[78] Smulders, P., 1990. The threat of information theft by reception of electromagnetic radiation from RS-232 cables. *Computers and Security*, 9(1), 53-58.

[79] SST, SST: Secure products for a more secure world. Available at: http://www.sst.ws/ [Accessed July 20, 2010].

[80] Stabernack, B., Wels, K. & Hubert, H., 2007. A Video Coprocessor for Mobile Multi Media Signal Processing. In *Proceedings of the IEEE International Symposium on Consumer Electronics, 2007.* pp. 1-6.

[81] Szczechowiak, P. et al., 2008. NanoECC: testing the limits of elliptic curve cryptography in sensor networks. In *Proceedings of the 5th European conference on Wireless sensor networks.* pp. 305-320.

[82] Tanaka, H., Takizawa, O. & Yamamura, A., 2005. Evaluation and Improvement of the Tempest Fonts. In *Information Security Applications.* pp. 457-469.

[83] The Guardian, 2010. Lebanese phone firm manager accused of spying for Israel. Available at: http://www.guardian.co.uk/world/2010/jun/30/lebanon-arrests-suspected-israeli-spy [Accessed July 9, 2010].

[84] The Register, India puts threatened BlackBerry ban on paper. Available at: http://www.theregister.co.uk/2010/08/17/india_rim_again/ [Accessed August 17, 2010].

[85] The Register, Security flaw creates Android, Palm Pre snoop risk. Available at: http://www.theregister.co.uk/2010/08/13/smartphone_security_bug/ [Accessed August 15, 2010].

[86] Thiele, E., Tempest for Eliza. Available at: http://www.erikyyy.de/tempest/ [Accessed July 20, 2010].

[87] Thorpe, J., Oorschot, P.C.V. & Somayaji, A., 2005. Pass-thoughts: authenticating with our minds. In *Proceedings of the 2005 workshop on New security paradigms.* pp. 45-56.

[88] Vuagnoux, M. & Pasini, S., 2009. Compromising Electromagnetic Emanations of Wired and Wireless Keyboards. In *Proceedings of the 18th USENIX Security Symposium.* pp. 1-16.

[89] Vuzix, Vuzix - View the Future Today. Available at: http://www.vuzix.com/home/index.html [Accessed July 4, 2010].

[90] Watanabe, T., Nagayoshi, H. & Sako, H., 2008. A Display Technique for Preventing Electromagnetic Eavesdropping Using Color Mixture Characteristic of Human Eyes. In *Information Hiding.* pp. 1-14.

[91] Wired, FBI Gets Evidence in Student Webcam Scandal | Threat Level. Available at: http://www.wired.com/threatlevel/2010/05/webcam-scandal-evidence/ [Accessed August 17, 2010].

[92] Wired, Forget the Drones: Executive Plane Now an Afghanistan Flying Spy | Danger Room. Available at: http://www.wired.com/dangerroom/2010/08/executive-plane-becomes-flying-spy-in-afghanistan/ [Accessed August 17, 2010].

[93] Wright, P., 1987. *Spycatcher: The Candid Autobiography of a Senior Intelligence Officer*, Viking Press.

[94] Wu, X. et al., 2009. Surveillance Robot Utilizing Video and Audio Information. *Journal of Intelligent and Robotic Systems*, 55(4), 403-421.

[95] Zhuang, L., Zhou, F. & Tygar, J.D., 2005. Keyboard acoustic emanations revisited. In *Proceedings of the 12th ACM conference on Computer and communications security*. pp. 373-382.

[96] Zouari, L. & Chollet, G., 2009. Efficient codebooks for fast and accurate low resource ASR systems. *Speech Communication*, 51(9), 732-743.